On the Quantum Complexity of Hawking Radiation

Daniel Harlow

Princeton University - PCTS

June 25, 2013

A Status Report on Firewalls

- Last year the previous speaker and some of his fellow miscreants presented an argument that one-sided black hole horizons are typically not smooth. (Almheiri/Marolf/Polchinski/Sully 2012)

A Status Report on Firewalls

- Last year the previous speaker and some of his fellow miscreants presented an argument that one-sided black hole horizons are typically not smooth. (Almheiri/Marolf/Polchinski/Sully 2012)
- Their argument, a version of which I'll briefly review below, begins with a very simple set of assumptions and has proven robust enough that in the ensuing time nobody has been able to really convincingly refute it.

A Status Report on Firewalls

- Last year the previous speaker and some of his fellow miscreants presented an argument that one-sided black hole horizons are typically not smooth. (Almheiri/Marolf/Polchinski/Sully 2012)
- Their argument, a version of which I'll briefly review below, begins with a very simple set of assumptions and has proven robust enough that in the ensuing time nobody has been able to really convincingly refute it.
- Despite $O(50)$ papers trying to!

A Status Report on Firewalls

- Last year the previous speaker and some of his fellow miscreants presented an argument that one-sided black hole horizons are typically not smooth. (Almheiri/Marolf/Polchinski/Sully 2012)
- Their argument, a version of which I'll briefly review below, begins with a very simple set of assumptions and has proven robust enough that in the ensuing time nobody has been able to really convincingly refute it.
- Despite $O(50)$ papers trying to!
- At this point it is clear that there remains something deep that we do not understand about black holes.

- When I say "convincingly refute", I do not just mean nitpicking with their assumptions or the technical details of their argument.

- When I say “convincingly refute”, I do not just mean nitpicking with their assumptions or the technical details of their argument.
- Although I will do both today!

- When I say “convincingly refute”, I do not just mean nitpicking with their assumptions or the technical details of their argument.
- Although I will do both today!
- What is really needed to kill the firewall is a simple quantum mechanical model of a one-sided black hole, formed from collapse, where it is simultaneously clear both how to describe the interior and the evaporation process, and in which both unitary and the validity of effective field theory for an infalling observer are apparent.

- When I say "convincingly refute", I do not just mean nitpicking with their assumptions or the technical details of their argument.
- Although I will do both today!
- What is really needed to kill the firewall is a simple quantum mechanical model of a one-sided black hole, formed from collapse, where it is simultaneously clear both how to describe the interior and the evaporation process, and in which both unitary and the validity of effective field theory for an infalling observer are apparent.
- Conversely to establish firewalls we'd need a model of how effective field theory breaks down, and an understanding of why it doesn't happen in other situations (like inflation).

- When I say “convincingly refute”, I do not just mean nitpicking with their assumptions or the technical details of their argument.
- Although I will do both today!
- What is really needed to kill the firewall is a simple quantum mechanical model of a one-sided black hole, formed from collapse, where it is simultaneously clear both how to describe the interior and the evaporation process, and in which both unitary and the validity of effective field theory for an infalling observer are apparent.
- Conversely to establish firewalls we’d need a model of how effective field theory breaks down, and an understanding of why it doesn’t happen in other situations (like inflation).
- I will not provide any such model today.

- When I say "convincingly refute", I do not just mean nitpicking with their assumptions or the technical details of their argument.
- Although I will do both today!
- What is really needed to kill the firewall is a simple quantum mechanical model of a one-sided black hole, formed from collapse, where it is simultaneously clear both how to describe the interior and the evaporation process, and in which both unitary and the validity of effective field theory for an infalling observer are apparent.
- Conversely to establish firewalls we'd need a model of how effective field theory breaks down, and an understanding of why it doesn't happen in other situations (like inflation).
- I will not provide any such model today.
- I will however introduce a set of ideas which may end up being an important ingredient in a firewall-free model. (Harlow/Hayden, 2013)

Operational Constraints

- My basic strategy is to study the *operational verifiability* of the contradiction at the heart of the AMPS argument.

Operational Constraints

- My basic strategy is to study the *operational verifiability* of the contradiction at the heart of the AMPS argument.
- This may seem misguided, after all isn't a contradiction unacceptable regardless of whether or not we can see it?

Operational Constraints

- My basic strategy is to study the *operational verifiability* of the contradiction at the heart of the AMPS argument.
- This may seem misguided, after all isn't a contradiction unacceptable regardless of whether or not we can see it?

Not necessarily!

Operational Constraints

- My basic strategy is to study the *operational verifiability* of the contradiction at the heart of the AMPS argument.
- This may seem misguided, after all isn't a contradiction unacceptable regardless of whether or not we can see it?

Not necessarily!

- Consider the uncertainty principle. Heisenberg argued that various practical restrictions prevent us from measuring both the position and the momentum of a particle with arbitrary precision. A classical physicist might say "who cares", I *know* that it has both a position and a momentum regardless of whether or not *you* can measure them!

Operational Constraints

- My basic strategy is to study the *operational verifiability* of the contradiction at the heart of the AMPS argument.
- This may seem misguided, after all isn't a contradiction unacceptable regardless of whether or not we can see it?

Not necessarily!

- Consider the uncertainty principle. Heisenberg argued that various practical restrictions prevent us from measuring both the position and the momentum of a particle with arbitrary precision. A classical physicist might say "who cares", I *know* that it has both a position and a momentum regardless of whether or not *you* can measure them!
- Obviously the classical physicist is wrong. The operational restrictions on measuring both the position and momentum enable the particle to behave in a way that would be contradictory with classical assumptions, but the problem is with those assumptions, not with the particle.

- A less successful example is S-matrix theory - this was a reaction against quantum fields as being “unobservable”, and an insistence that theories should refer only to things that are actually measureable.

- A less successful example is S-matrix theory - this was a reaction against quantum fields as being "unobservable", and an insistence that theories should refer only to things that are actually measureable.
- This turned out to be misguided; the fields are essential in understanding the renormalization group, instantons, etc.

- A less successful example is S-matrix theory - this was a reaction against quantum fields as being “unobservable”, and an insistence that theories should refer only to things that are actually measureable.
- This turned out to be misguided; the fields are essential in understanding the renormalization group, instantons, etc.
- Of course S-matrix theory did lead to some minorly interesting ideas anyway...

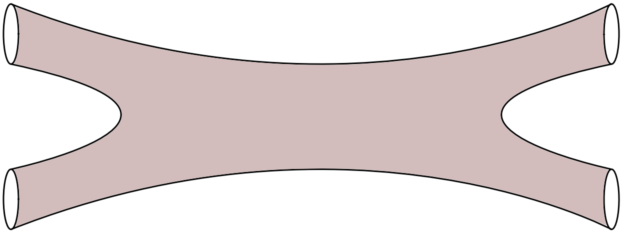

Another example, closer to the topic at hand, is the old discussion of cloning and unitarity in black hole evaporation:

Another example, closer to the topic at hand, is the old discussion of cloning and unitarity in black hole evaporation:

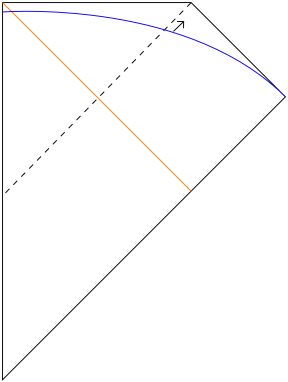

Another example, closer to the topic at hand, is the old discussion of cloning and unitarity in black hole evaporation:

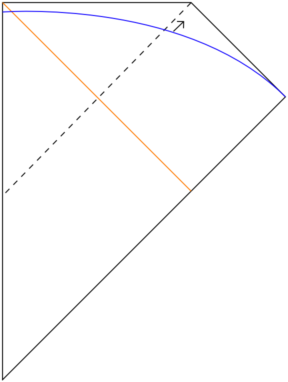

- Here the operational restrictions come from *causality*.

Another example, closer to the topic at hand, is the old discussion of cloning and unitarity in black hole evaporation:

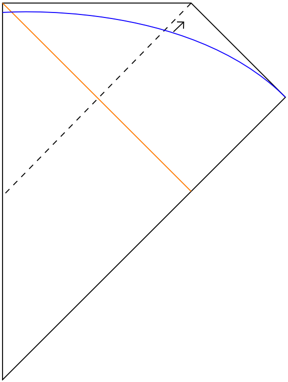

- Here the operational restrictions come from *causality*.
- The idea of *complementarity* was that the infalling shell and the Hawking radiation are not really distinct degrees of freedom in the same Hilbert space; this is ok because nobody can see both! Just like position and momentum...

Firewalls!

The AMPS argument is a reformulation of the information loss problem in which all of the moving parts are causally accessible to a single observer - apparently eliminating the possibility of the problem being resolved by complementarity.

Firewalls!

The AMPS argument is a reformulation of the information loss problem in which all of the moving parts are causally accessible to a single observer - apparently eliminating the possibility of the problem being resolved by complementarity.

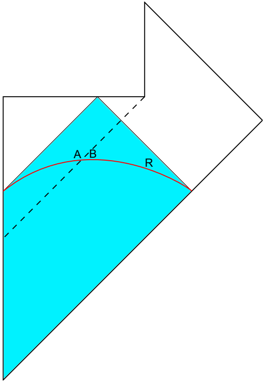

- "A" are some right-moving Rindler modes just inside the horizon.

Firewalls!

The AMPS argument is a reformulation of the information loss problem in which all of the moving parts are causally accessible to a single observer - apparently eliminating the possibility of the problem being resolved by complementarity.

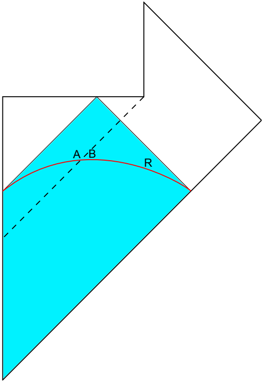

- "A" are some right-moving Rindler modes just inside the horizon.
- "B" are some right-moving Rindler modes just outside the horizon.

Firewalls!

The AMPS argument is a reformulation of the information loss problem in which all of the moving parts are causally accessible to a single observer - apparently eliminating the possibility of the problem being resolved by complementarity.

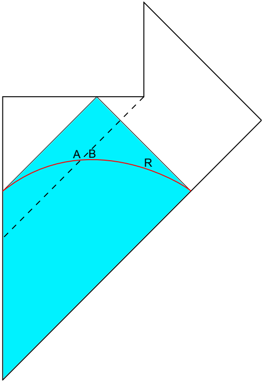

- “A” are some right-moving Rindler modes just inside the horizon.
- “B” are some right-moving Rindler modes just outside the horizon.
- “R” is the Hawking radiation that has been emitted so far.

The goal is then to come up with some quantum state on the red slice that is consistent with both unitarity and a smooth horizon.

The goal is then to come up with some quantum state on the red slice that is consistent with both unitarity and a smooth horizon.

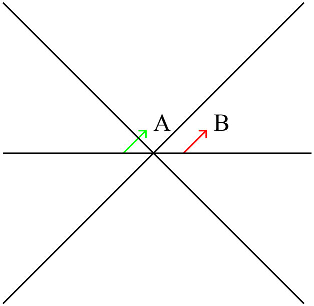

From Rindler space we know that a smooth vacuum requires A and B to be quite entangled, which I denote as $A \longleftrightarrow B$.

The goal is then to come up with some quantum state on the red slice that is consistent with both unitarity and a smooth horizon.

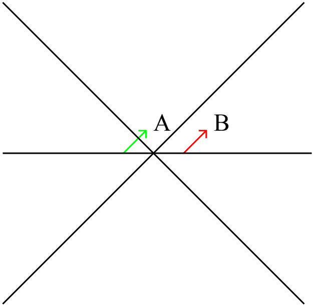

From Rindler space we know that a smooth vacuum requires A and B to be quite entangled, which I denote as $A \longleftrightarrow B$.
More quantitatively, we can write the Minkowski vacuum as:

$$|\Omega\rangle = \frac{1}{Z}\sum_i e^{-\frac{\beta\omega_i}{2}}|i\rangle_L|i\rangle_R. \quad (1)$$

- The AMPS paradox arises because, once the black hole is "old" in the sense of having radiated away more than half of its entropy, unitarity $+$ scrambling also require B to be entangled with the radiation.

- The AMPS paradox arises because, once the black hole is "old" in the sense of having radiated away more than half of its entropy, unitarity + scrambling also require B to be entangled with the radiation.
- More carefully, they require B to be entangled with some *subfactor* of R, usually denoted R_B, whose Hilbert space dimensionality is equal to that of B.

- The AMPS paradox arises because, once the black hole is "old" in the sense of having radiated away more than half of its entropy, unitarity + scrambling also require B to be entangled with the radiation.
- More carefully, they require B to be entangled with some *subfactor* of R, usually denoted R_B, whose Hilbert space dimensionality is equal to that of B.
- R_B is often called the *purification* of B, in the sense that it is the smallest subfactor for which $S_{BR_B} = 0$.

- The AMPS paradox arises because, once the black hole is "old" in the sense of having radiated away more than half of its entropy, unitarity + scrambling also require B to be entangled with the radiation.
- More carefully, they require B to be entangled with some *subfactor* of R, usually denoted R_B, whose Hilbert space dimensionality is equal to that of B.
- R_B is often called the *purification* of B, in the sense that it is the smallest subfactor for which $S_{BR_B} = 0$.
- This is a problem from the point of view of strong subadditivity:

$$S_{ABR_B} + S_B \leq S_{AB} + S_{BR_B} \tag{2}$$

- This inequality, along with $S_{BR_B} = 0$, implies that the mutual information $S_A + S_B - S_{AB}$ is zero. This is completely inconsistent with $A \longleftrightarrow B$!

The Overlap Rule

You might object that since Alice falls into the black hole and hits a singularity she does not need to be constrained by unitarity.

The Overlap Rule

You might object that since Alice falls into the black hole and hits a singularity she does not need to be constrained by unitarity.
But consider Charlie:

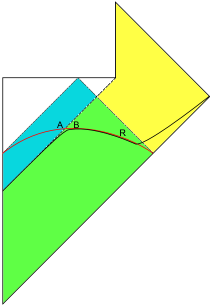

The Overlap Rule

You might object that since Alice falls into the black hole and hits a singularity she does not need to be constrained by unitarity.
But consider Charlie:

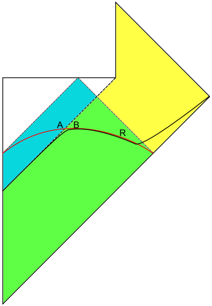

Since B and R are "obviously" accessible to both Alice and Charlie we apparently need

$$\rho_{BR}^{\{Alice\}} = \rho_{BR}^{\{Charlie\}}. \tag{3}$$

A New Kind of Complementarity?

In the remainder of this talk I will argue that in fact the embedding of R_B into the radiation is so sophisticated that Alice is unable to verify the entanglement of B and R_B and still have time to jump into the black hole. This may allow a new kind of complementarity:

A New Kind of Complementarity?

In the remainder of this talk I will argue that in fact the embedding of R_B into the radiation is so sophisticated that Alice is unable to verify the entanglement of B and R_B and still have time to jump into the black hole. This may allow a new kind of complementarity:

- Charlie sees $B \longleftrightarrow R_B$ as demanded by unitarity, but cannot access A since it is behind the horizon.

A New Kind of Complementarity?

In the remainder of this talk I will argue that in fact the embedding of R_B into the radiation is so sophisticated that Alice is unable to verify the entanglement of B and R_B and still have time to jump into the black hole. This may allow a new kind of complementarity:

- Charlie sees $B \longleftrightarrow R_B$ as demanded by unitarity, but cannot access A since it is behind the horizon.
- Alice sees $A \longleftrightarrow B$ as required for smooth infall, but cannot access R_B since it is too computationally difficult.

A New Kind of Complementarity?

In the remainder of this talk I will argue that in fact the embedding of R_B into the radiation is so sophisticated that Alice is unable to verify the entanglement of B and R_B and still have time to jump into the black hole. This may allow a new kind of complementarity:

- Charlie sees $B \longleftrightarrow R_B$ as demanded by unitarity, but cannot access A since it is behind the horizon.
- Alice sees $A \longleftrightarrow B$ as required for smooth infall, but cannot access R_B since it is too computationally difficult.

I interpret this as suggesting that perhaps the AMPS assumptions about the structure of the Hilbert space, like thinking that a particle has both a position and a momentum, are too restrictive.

Alice's Task

To make this argument, I need to be more precise about what exactly Alice needs to do to verify the entanglement of B and R_B.

Alice's Task

To make this argument, I need to be more precise about what exactly Alice needs to do to verify the entanglement of B and R_B.
It will be convenient to set things up in Charlie's description, which by the overlap rule is equivalent.

Alice's Task

To make this argument, I need to be more precise about what exactly Alice needs to do to verify the entanglement of B and R_B.
It will be convenient to set things up in Charlie's description, which by the overlap rule is equivalent.
In Charlie's unitary theory we can think of the state of the system as being a pure state in a Hilbert space that, following AMPS, I'll take to factorize as

$$\mathcal{H} = \mathcal{H}_H \otimes \mathcal{H}_B \otimes \mathcal{H}_R. \quad (4)$$

Alice's Task

To make this argument, I need to be more precise about what exactly Alice needs to do to verify the entanglement of B and R_B.
It will be convenient to set things up in Charlie's description, which by the overlap rule is equivalent.
In Charlie's unitary theory we can think of the state of the system as being a pure state in a Hilbert space that, following AMPS, I'll take to factorize as

$$\mathcal{H} = \mathcal{H}_H \otimes \mathcal{H}_B \otimes \mathcal{H}_R. \tag{4}$$

Here H is the stretched horizon degrees of freedom, B is quantum field theory degrees of freedom inside the angular momentum barrier near $r \approx 3GM$, and R is the radiation field outside of the barrier.

From now on I will model all subfactors as being made of qubits, with n qubits for R, m qubits for H, and k qubits for B. To a first approximation you can think of all of these as being of order the entropy of the black hole, although for the black hole to be "old" we need $n > k + m$.

From now on I will model all subfactors as being made of qubits, with n qubits for R, m qubits for H, and k qubits for B. To a first approximation you can think of all of these as being of order the entropy of the black hole, although for the black hole to be "old" we need $n > k + m$.
For an "old" black hole we expect the state to have the form

$$|\Psi\rangle = \frac{1}{\sqrt{|B||H|}} \sum_{b,h} |b\rangle_B |h\rangle_H U_R |bh0\rangle_R, \tag{5}$$

where more explicitly

$$\begin{aligned} |b\rangle_B &= |b_1 \dots b_k\rangle_B \\ |h\rangle_H &= |h_1 \dots h_m\rangle_H \\ |bh0\rangle_R &= |b_1 \dots b_k h_1 \dots h_k 0 \dots 0\rangle_R . \end{aligned} \tag{6}$$

From now on I will model all subfactors as being made of qubits, with n qubits for R, m qubits for H, and k qubits for B. To a first approximation you can think of all of these as being of order the entropy of the black hole, although for the black hole to be "old" we need $n > k + m$.
For an "old" black hole we expect the state to have the form

$$|\Psi\rangle = \frac{1}{\sqrt{|B||H|}} \sum_{b,h} |b\rangle_B |h\rangle_H U_R |bh0\rangle_R, \tag{5}$$

where more explicitly

$$\begin{aligned} |b\rangle_B &= |b_1 \dots b_k\rangle_B \\ |h\rangle_H &= |h_1 \dots h_m\rangle_H \\ |bh0\rangle_R &= |b_1 \dots b_k h_1 \dots h_k 0 \dots 0\rangle_R. \end{aligned} \tag{6}$$

Here I am interpreting the bases for B and R as being simple from the point of view of local measurements, so U_R is the complicated unitary transformation on the radiation that relates this local basis to a basis (the Schmidt basis) where the entanglement is manifest.

To verify the entanglement Alice simply needs to act on the system with $U_R^\dagger$.

- Acting with your favorite unitary transformation on some system is called doing a *quantum computation*.

To verify the entanglement Alice simply needs to act on the system with $U_R^\dagger$.

- Acting with your favorite unitary transformation on some system is called doing a *quantum computation*.
- The study of how hard it is to do this is called *quantum complexity theory*. The basic question is how the amount of memory and time required scales with n, the number of input bits.

To verify the entanglement Alice simply needs to act on the system with $U_R^\dagger$.

- Acting with your favorite unitary transformation on some system is called doing a *quantum computation*.
- The study of how hard it is to do this is called *quantum complexity theory*. The basic question is how the amount of memory and time required scales with n, the number of input bits.

For a Schwarzschild black hole the evaporation time scales like $n^{\frac{3}{2}}$, so from a complexity-theoretic point of view Alice needs to be able to compute quite fast in order to be able to finish before the black hole evaporates.

To verify the entanglement Alice simply needs to act on the system with $U_R^\dagger$.

- Acting with your favorite unitary transformation on some system is called doing a *quantum computation*.
- The study of how hard it is to do this is called *quantum complexity theory*. The basic question is how the amount of memory and time required scales with n, the number of input bits.

For a Schwarzschild black hole the evaporation time scales like $n^{\frac{3}{2}}$, so from a complexity-theoretic point of view Alice needs to be able to compute quite fast in order to be able to finish before the black hole evaporates. For comparison, multiplying two $n \times n$ matrices with $O(1)$ coefficients naively requires time n^3, although it has been optimized so far to $n^{2.3}$. By contrast in quantum mechanics we are typically multiplying $2^n \times 2^n$ matrices!

To verify the entanglement Alice simply needs to act on the system with $U_R^\dagger$.

- Acting with your favorite unitary transformation on some system is called doing a *quantum computation*.
- The study of how hard it is to do this is called *quantum complexity theory*. The basic question is how the amount of memory and time required scales with n, the number of input bits.

For a Schwarzschild black hole the evaporation time scales like $n^{\frac{3}{2}}$, so from a complexity-theoretic point of view Alice needs to be able to compute quite fast in order to be able to finish before the black hole evaporates. For comparison, multiplying two $n \times n$ matrices with $O(1)$ coefficients naively requires time n^3, although it has been optimized so far to $n^{2.3}$. By contrast in quantum mechanics we are typically multiplying $2^n \times 2^n$ matrices!
Indeed Patrick and I argued that verifying the entanglement between B and R_B typically takes a time of order $2^{\# n}$.

The Quantum Circuit Model

We need some model of a quantum computer to assess how hard it is to make a unitary transformation.

The Quantum Circuit Model

We need some model of a quantum computer to assess how hard it is to make a unitary transformation. The standard model in the literature is to build up elements of $U(2^n)$ by acting with elements of $U(4)$, called *gates*, on the qubits two at a time:

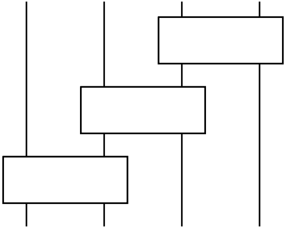

The Quantum Circuit Model

We need some model of a quantum computer to assess how hard it is to make a unitary transformation. The standard model in the literature is to build up elements of $U(2^n)$ by acting with elements of $U(4)$, called *gates*, on the qubits two at a time:

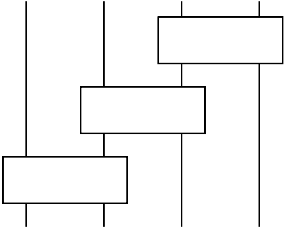

It is a theorem that even with a single sufficiently generic type of gate, we can approximate an arbitrary element of $U(2^n)$ to arbitrary accuracy by applying this gate to various pairs of qubits in succession!

A (finite) sequence of gates is called a *quantum circuit* and the number of gates is called its *size*.

A (finite) sequence of gates is called a *quantum circuit* and the number of gates is called its *size*.

The size is a good measure of the complexity of the circuit. At least if we apply the gates in sucession it is directly proportional to the time needed for the computation.

A (finite) sequence of gates is called a *quantum circuit* and the number of gates is called its *size*.

The size is a good measure of the complexity of the circuit. At least if we apply the gates in sucession it is directly proportional to the time needed for the computation.

How many gates are needed to approximate a generic unitary to accuracy ϵ? Roughly

$$T \approx 2^{2n} \log \frac{1}{\epsilon}. \tag{7}$$

A (finite) sequence of gates is called a *quantum circuit* and the number of gates is called its *size*.

The size is a good measure of the complexity of the circuit. At least if we apply the gates in sucession it is directly proportional to the time needed for the computation.

How many gates are needed to approximate a generic unitary to accuracy ϵ? Roughly

$$T \approx 2^{2n} \log \frac{1}{\epsilon}. \tag{7}$$

To see this one observes that the number of circuits of size T is about

$$\left(2\binom{n}{2}\right)^T \approx n^{2T}, \tag{8}$$

while the number of balls needed to cover $U(2^n)$ is of order

$$\left(\frac{1}{\epsilon}\right)^{2^{2n}}. \tag{9}$$

So indeed we see that the time required is exponential in the entropy for almost all possible U_R's. Some comments:

So indeed we see that the time required is exponential in the entropy for almost all possible U_R's. Some comments:

- Doing more than one gate at a time will not help, since we'd need to do exponentially many at once, and then just communicating from one part of the computer to the other would take exponential time.

So indeed we see that the time required is exponential in the entropy for almost all possible U_R's. Some comments:

- Doing more than one gate at a time will not help, since we'd need to do exponentially many at once, and then just communicating from one part of the computer to the other would take exponential time.
- Variations of the model like increasing the number of gates or using higher spin fundamental objects do nothing to this estimate.

So indeed we see that the time required is exponential in the entropy for almost all possible U_R's. Some comments:

- Doing more than one gate at a time will not help, since we'd need to do exponentially many at once, and then just communicating from one part of the computer to the other would take exponential time.
- Variations of the model like increasing the number of gates or using higher spin fundamental objects do nothing to this estimate.
- This counting is the fundamental reason that quantum computers still cannot solve really hard problems in polynomial time.

Why Not Time Reverse?

It may seem that since the black hole made the radiation in polynomial time, we should be able to decode it in comparable time by some kind of time reversal.

Why Not Time Reverse?

It may seem that since the black hole made the radiation in polynomial time, we should be able to decode it in comparable time by some kind of time reversal.

To address this, we need to understand a little better what the black hole does.

Why Not Time Reverse?

It may seem that since the black hole made the radiation in polynomial time, we should be able to decode it in comparable time by some kind of time reversal.

To address this, we need to understand a little better what the black hole does.

Here is a model black hole dynamics:

$$|\Psi\rangle = U_{dyn}|i\rangle = \frac{1}{\sqrt{|B||H|}} \sum_{b,h} |b\rangle_B |h\rangle_H U_R(i)|bh0\rangle_R \tag{10}$$

Why Not Time Reverse?

It may seem that since the black hole made the radiation in polynomial time, we should be able to decode it in comparable time by some kind of time reversal.

To address this, we need to understand a little better what the black hole does.

Here is a model black hole dynamics:

$$|\Psi\rangle = U_{dyn}|i\rangle = \frac{1}{\sqrt{|B||H|}} \sum_{b,h} |b\rangle_B |h\rangle_H U_R(i)|bh0\rangle_R \tag{10}$$

Here U_{dyn} is the time evolution operator, which for various reasons we should expect has a polynomial-sized circuit.

Why Not Time Reverse?

It may seem that since the black hole made the radiation in polynomial time, we should be able to decode it in comparable time by some kind of time reversal.

To address this, we need to understand a little better what the black hole does.

Here is a model black hole dynamics:

$$|\Psi\rangle = U_{dyn}|i\rangle = \frac{1}{\sqrt{|B||H|}} \sum_{b,h} |b\rangle_B |h\rangle_H U_R(i)|bh0\rangle_R \tag{10}$$

Here U_{dyn} is the time evolution operator, which for various reasons we should expect has a polynomial-sized circuit.

Using simple manipulations we can construct from U_{dyn} a polynomial sized circuit $\tilde{U}_R^\dagger$ with the property that it acts on $|\Psi\rangle$ as $U_R^\dagger$.

Why Not Time Reverse?

It may seem that since the black hole made the radiation in polynomial time, we should be able to decode it in comparable time by some kind of time reversal.

To address this, we need to understand a little better what the black hole does.

Here is a model black hole dynamics:

$$|\Psi\rangle = U_{dyn}|i\rangle = \frac{1}{\sqrt{|B||H|}} \sum_{b,h} |b\rangle_B |h\rangle_H U_R(i)|bh0\rangle_R \tag{10}$$

Here U_{dyn} is the time evolution operator, which for various reasons we should expect has a polynomial-sized circuit.

Using simple manipulations we can construct from U_{dyn} a polynomial sized circuit $\tilde{U}_R^\dagger$ with the property that it acts on $|\Psi\rangle$ as $U_R^\dagger$.

But its gates act nontrivially on B and H, so it is useless for Alice! In other words, inverting a unitary is much harder when you don't have access to the whole system...

In fact since for each initial state $|i\rangle$ requires its own $U_R(i)$, there is a simple counting argument that there are far too many initial states for more than a vanishingly small fraction to even have a chance of having a $U_R(i)$ of polynomial size.

In fact since for each initial state $|i\rangle$ requires its own $U_R(i)$, there is a simple counting argument that there are far too many initial states for more than a vanishingly small fraction to even have a chance of having a $U_R(i)$ of polynomial size.
This is a rigorous proof that the decoding will take exponential time for generic initial states!

In fact since for each initial state $|i\rangle$ requires its own $U_R(i)$, there is a simple counting argument that there are far too many initial states for more than a vanishingly small fraction to even have a chance of having a $U_R(i)$ of polynomial size.

This is a rigorous proof that the decoding will take exponential time for generic initial states!

One might still wonder if there is a basis of initial states where the decoding is easy; I won't have time to discuss this here, but using results from the computational complexity of quantum error correction we were able to argue that this is very unlikely.

In fact since for each initial state $|i\rangle$ requires its own $U_R(i)$, there is a simple counting argument that there are far too many initial states for more than a vanishingly small fraction to even have a chance of having a $U_R(i)$ of polynomial size.

This is a rigorous proof that the decoding will take exponential time for generic initial states!

One might still wonder if there is a basis of initial states where the decoding is easy; I won't have time to discuss this here, but using results from the computational complexity of quantum error correction we were able to argue that this is very unlikely.

You could also ask if some sort of exotic computer, perhaps a carefully constructed black hole, could help speed up the computation. Without knowing the laws of quantum gravity we couldn't be sure, but we were able to give an argument that the existence of such a "hypercomputer" is extremely unlikely without the theory having some sort of special structure which is specifically designed to allow this.

Other Black Holes

For Schwarzschild black holes it seems pretty clear that the power-law evaporation time is obliterated by this exponential. One can try to beat it by considering other types of black holes.

Other Black Holes

For Schwarzschild black holes it seems pretty clear that the power-law evaporation time is obliterated by this exponential. One can try to beat it by considering other types of black holes.

Simple things like adding charge, angular momentum, putting in AdS don't seem to help.

Other Black Holes

For Schwarzschild black holes it seems pretty clear that the power-law evaporation time is obliterated by this exponential. One can try to beat it by considering other types of black holes.

Simple things like adding charge, angular momentum, putting in AdS don't seem to help.

The most dangerous candidate seems to be taking a large black hole in AdS and putting it exponentially far down a long but finite throat that asymptotes to 10D Minkowski space:

$$ds^2 = \frac{1}{\sqrt{1+\left(\frac{R}{r}\right)^4}}\left(-dt^2+dx^2\right)+\sqrt{1+\left(\frac{R}{r}\right)^4}\left(dr^2+r^2d\Omega^2\right). \quad (11)$$

Other Black Holes

For Schwarzschild black holes it seems pretty clear that the power-law evaporation time is obliterated by this exponential. One can try to beat it by considering other types of black holes.

Simple things like adding charge, angular momentum, putting in AdS don't seem to help.

The most dangerous candidate seems to be taking a large black hole in AdS and putting it exponentially far down a long but finite throat that asymptotes to 10D Minkowski space:

$$ds^2 = \frac{1}{\sqrt{1+\left(\frac{R}{r}\right)^4}}\left(-dt^2+dx^2\right)+\sqrt{1+\left(\frac{R}{r}\right)^4}\left(dr^2+r^2d\Omega^2\right). \quad (11)$$

The black hole then decays very slowly by losing energy out to the Minkowski region, and one can imagine "outsourcing" the computation up the throat to the 10D region.

- This construction introduces a new problem however in that the barrier makes it very difficult to send signals reliably down the throat without destroying the black hole, and so far the timescales seem to work out in such a way that it remains impossible for Alice to verify the entanglement.

- This construction introduces a new problem however in that the barrier makes it very difficult to send signals reliably down the throat without destroying the black hole, and so far the timescales seem to work out in such a way that it remains impossible for Alice to verify the entanglement.
- Recently AMPSS have tried to evade these “joining pains” by replacing the Asymptotic Minkowski region by some arbitrary system with no spacetime interpretation. This construction raises confusing interpretational questions about quantum gravity, and in particular it seems to me that some of their discussion is inconsistent with my understanding of the eternal two-sided AdS black hole. (See Juan’s talk)

- This construction introduces a new problem however in that the barrier makes it very difficult to send signals reliably down the throat without destroying the black hole, and so far the timescales seem to work out in such a way that it remains impossible for Alice to verify the entanglement.
- Recently AMPSS have tried to evade these "joining pains" by replacing the Asymptotic Minkowski region by some arbitrary system with no spacetime interpretation. This construction raises confusing interpretational questions about quantum gravity, and in particular it seems to me that some of their discussion is inconsistent with my understanding of the eternal two-sided AdS black hole. (See Juan's talk)
- I think that for now it is best to focus on thought experiments done strictly by observers living in the bulk with well-defined low-energy initial conditions; otherwise I think we can't really be too sure what to expect.

Conclusion

In the end, there is still much more to understand.

Conclusion

In the end, there is still much more to understand.
As I said in the beginning, these ideas have not yet been implemented into a complete model of black hole physics that takes advantage of them to circumvent the AMPS argument. As such I am not yet satisfied.

Conclusion

In the end, there is still much more to understand.

As I said in the beginning, these ideas have not yet been implemented into a complete model of black hole physics that takes advantage of them to circumvent the AMPS argument. As such I am not yet satisfied.

Thanks for listening! I'll leave you with a picture I found online of a black hole formation and evaporation

Conclusion

In the end, there is still much more to understand.
As I said in the beginning, these ideas have not yet been implemented into a complete model of black hole physics that takes advantage of them to circumvent the AMPS argument. As such I am not yet satisfied.
Thanks for listening! I'll leave you with a picture I found online of a black hole formation and evaporation