Measurement-Induced Phase Transitions

Matthew Thibodeau

Abstract

When some unitary quantum circuits are punctuated by periodic measurements, distinct phases can emerge, where the order parameter is related to the entanglement properties of the circuits. This term paper describes recent progress in the identification and characterization of these phases. We provide a selection of circuit models that exhibit this phenomenon, detail the analytic and numerical arguments that show the emergence of high- and low-entanglement phases, and then describe some important properties of the circuits and states that can emerge, including topological protection and quantum error correction capabilities.

I. INTRODUCTION

The looming spectre of *quantum computation* has commandeered the talent and imagination of a generation of information theorists. Part of their effort has been to understand the fundamental resources that enable quantum computation to advance past its classical counterpart; foremost among them is entanglement between spatially separated degrees of freedom. To that end, the entangling properties of different quantum circuits, which are the basic units of quantum computation, are interesting not only in their own right but also through their potential applications to quantum algorithms, including, as we will see, the ever-elusive goal of error correction.

According to the postulates of quantum mechanics, a system evolves in time under the action of a unitary operator U until it is "measured," at which point it is projected into one eigenstate of the measurement operator P. Entanglement in the system can only be generated by U and is generally destroyed by the action of P. A natural question to ask is, then: how do these effects behave when placed in competition with one another? We will see that there is a delicate interplay between them, with a bona-fide phase transition between a parameter regime where U wins unambiguously and another where P dominates.

We'll begin in earnest by describing the simplest case of the entanglement transition on a 1D array of qubits with very simple discrete-time dynamics and projective measurements of σ_z, which can be mapped onto a problem of classical percolation. Here we introduce the concepts of *entangling* and *disentangling* phases, which are the basic results of an entanglement phase transition. Next, we broaden the setting to include different types of measurements and examine the possible interesting features of the phases, including topological protection. Finally, we close with a discussion of the interpretation of the problem from the point of view of quantum error correction, where we will find that both phases have some ability to protect quantum information from destruction by projective measurements, and a perspective on future research directions and $D > 1$ continuous time local models.

II. ENTANGLEMENT TRANSITIONS IN 1D

Entanglement phases are naturally suited for study on a quantum computer, which is the setting we will focus on in this paper. In particular, we'll start with an array of qubits in one dimension which are entangled by unitary quantum circuits and then dis-entangled by projective measurements of their spin, a problem that was treated extensively in [1].

A. Block Circuits and Entropy

We'll specialize to a particularly simple type of random circuit, referred to as a block circuit, illustrated in Figure 1. The array of n qubits lies on the horizontal axis and time runs vertically. This is a "discrete-time" quantum circuit, because each integer timestep is marked by the application of a layer of two-qubit unitary operators U chosen uniformly at random (with respect to the Haar measure), represented by boxes, followed by some randomly chosen measurements of σ_z, represented by dots. The only free parameter in this model is the rate p at which measurements are conducted: During each timestep, each qubit has its σ_z measured with probability p.

The unitary operators tend to entangle the two qubits that they touch, and because of

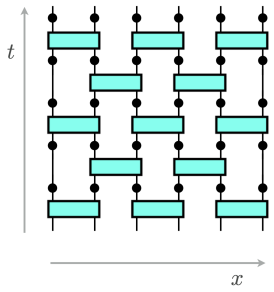

FIG. 1: A block circuit, showing unitaries (teal blocks) interspersed with projectors (black dots). Adapted from [1].

their alternating structure, after a long enough time they connect every pair of qubits; thus, they tend to drive the state to what's known as *volume-law* entanglement, where the amount of entanglement (the measurement of which we will discuss next) between a subsystem A and its complement $\bar{A}$ scales like the number of qubits in A, which we identify as its volume. On the other hand, the projective measurements tend to destroy entanglement, because after qubit j is measured to be in state $|x\rangle_j$, the system is certainly in a product state $|x\rangle_j \otimes |\Psi\rangle_{\hat{j}}$, where $|\Psi\rangle_{\hat{j}}$ is the state of the system excluding qubit j.

We say that a system has *area-law* entanglement if the entanglement of subsystem A scales with the area of the interface between A and $\bar{A}$. In 1D, this a constant because the interface is a point if A is connected, so in 1D, product states (with zero entanglement entropy) satisfy the area-law. Thus, projective measurements drive the state to area-law, and the tension between the competing effects of the unitaries and projectors is what causes the phase transition.

Entanglement between two subsystems of a quantum state ρ, which we represent by its density matrix, is somewhat difficult to quantify. The classic measures are the Renyi entropies: given a subsystem A of the full set of qubits, and $n \geq 0$, the entropy of ρ with respect to A is

$$S_n(\rho) = \frac{1}{1-n} \operatorname{Tr}_A \log_2 \rho_A^n \qquad S_0(\rho) = \log_2 N \tag{II.1}$$

where $\rho_A = \operatorname{Tr}_{\bar{A}} \rho$ is the reduced density matrix. In the limit where $n \to 0$, the trace counts the number N of nonzero eigenvalues of ρ_A. We will focus on S_0 because it has a nice interpretation in terms of "bond-cutting" of the random unitaries that are present in our circuit.

B. Minimal Cuts, Percolation, and the Phase Transition

How can we easily calculate S_0? In the case where ρ is unentangled, ρ_A is a pure state and thus has only one eigenvalue (it is a projector), so $S_0(\text{unentangled state}) = \log_2 1 = 0$ vanishes. It turns out that there is an extremely intuitive way to extend this calculation to the case where ρ begins as a product state and is then entangled by a block-like circuit such as the one shown in Figure 1: S_0 is exactly equal to the minimum number of links, known as the *minimal cut*, that one needs to cut to be able to separate off the subsystem A. The

process is shown in Figure 2, where to cut off the leftmost k spins one fixes the endpoint of the cut between spins k and $k+1$ and then draws the cut downward. This agrees with the requirement that $S_0 = 0$ for a product state, because there are no links that connect A and $\bar{A}$ in this case. In general, one must determine the path through the circuit that minimizes the number of bonds cut.

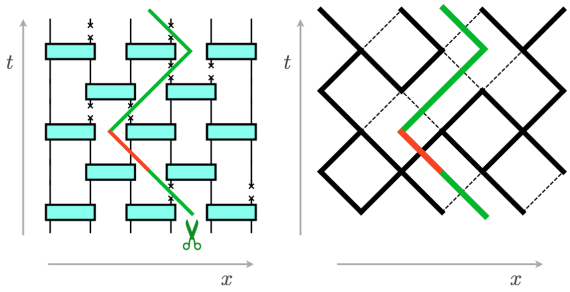

FIG. 2: On the left, part of a minimal cut in a block circuit, with some bonds broken by projections; only the red line contributes to the entropy. On the right, the associated lattice L of bonds. Adapted from [1].

This prescription is easily extended to the case where projective measurements are included. If a projective measurement occurs on the bond between two of the unitaries, we can regard that bond as being "broken" and thus not contributing to the count of our minimal cut. Intuitively, this works because even if the unitary entangles the two inputs, a projective measurement will yield a product state with, as discussed, no entropy across the cut bond. This is illustrated in Figure 2, which shows part of a minimal cut traversing some broken bonds.

With the given spatial form of the circuit, where the unitaries form a square lattice L tilted on its side as shown on the right in Figure 2, we have reduced the calculation of S_0 to finding the number of filled-in edges of the lattice that one must cut to travel between two points on the dual lattice L^* (made of the points in the center of each of the squares of the original lattice). This is a classical problem, known as *percolation*, with a known algorithmic solution and known asymptotic behavior of S_0 (in the discrete time coordinate t) as a function of p.

We are now ready to understand the phase transition, where the order parameter is the asymptotic limit of $S_0(t)$ between the two halves of the 1D system of qubits (i.e., A is equal to the left half of the system). There are three broad options for the scaling of S_0 as t becomes very large, which we can see by examining this percolation problem. These are illustrated in Figure 3.

1. $p \simeq 0$: *$S_0(t)$ grows linearly with t.* First, suppose that p is very small, so that nearly none of the edges of the lattice L are missing. Then nearly all steps through the lattice will incur a unit cost, so that the minimal cut is nearly a straight vertical line through the dual lattice L^*. The distance is approximately t, and the entropy is approximately $S_0 \simeq (1-p)t$. Note that for a finite system of size L, eventually S_0 must stop growing and will reach a maximum value proportional to the volume of A, which is $L/2$.

2. $p \simeq 1$: *$S_0(t)$ saturates at a small value, then stops growing.* Now, suppose that p is quite large, so that nearly none of the edges of L are filled in. Then there are

small "islands," regions of contiguous filled-in edges, but crucially, these islands are disconnected from one another. The minimal cut can wind between these islands, incurring no cost except if it needs to initially jump out of one. Thus, $S_0(t)$ does not grow with t (but may obtain a small value).

3. $p \simeq p_c$: *$S_0(t)$ grows logarithmically with t.* This is the case where p is near the transition point p_c between the two situations described above. The problem now possesses scale-invariance: there are now both islands of edges, as well as voids of missing edges, of all sizes. The minimal cut will generically make its way from a small void up through larger voids until it reaches one that is of size t, at which point it stops incurring an entropy cost. If each void was, say, double the size of the prior one, then there would be $\log_2 t$ jumps between voids, for a total cost proportional to $\log_2 t$.

For the square lattice, $p_c = \frac{1}{2}$.

These heuristic pictures are given a rigorous mathematical treatment in [2]. In our case, we can see that the first phase, where $S_0(t)$ grows linearly, corresponds to the *entangling phase* of the circuit, since it describes a situation in which the initial product state is ultimately endowed with volume-law entropy. The second phase is similarly referred to as a *disentangling phase*, because even if the state starts with volume-law, eventually the entropy will become a small constant. The marginal case in point three describes the transition between these two phases.

C. Experiments: Numerical Evidence for $n \geq 0$

The above is a solid theoretical argument for the existence of a phase transition in S_n for $n = 0$ in a 1D system, and it has been supported by numerical experiments. First, Ref [1] numerically simulated the percolation network and calculated both the growth rate of S_0 at short times as well as its asymptotic value as functions of p. These are shown in Figure 4, where they chose a maximum time of $t = 4L$, where L is the lattice width, to recover the asymptotic case. The results broadly agree with the theoretical predictions, and show the disappearance of S_0 after $p_c = 1/2$ as well as the linear growth at short times.

Conceivably, instead of relying on the results from percolation, the block circuits could be implemented directly on a real quantum computer, but no quantum computer currently exists that can run deep (that is, large t) circuits with high fidelity. However, we have an even better option: the circuits in question can be simulated directly on a classical computer using tensor network methods. Unlike the physical case, where the details of the state (e.g. the basis state amplitudes) cannot be directly determined, on a classical computer the entire state is recorded in memory and can be manipulated directly.

In particular, all of the Renyi entropies of the states, including S_0, can be calculated. However, S_0 is not a commonly used measure of entropy because it is sensitive to infinitesimally small perturbations away from the zero-entropy product state case. More often, the Von Neumann entropy $S_1 = -\operatorname{Tr}_A \rho_A \log \rho_A$ is preferred because it does not suffer from this problem. Direct simulation of the circuits is computationally expensive because of the exponential growth of Hilbert space with L, but Ref [1] also performed these calculations. The results are shown in Figure 5. Although the critical rate is different, with $p_c \simeq 0.25$, the data are broadly similar to the S_0 case – in particular, the figure shows that for $p > p_c$ the asymptotic limit of S_1 is both small and independent of L.

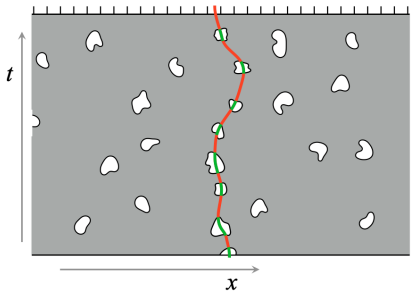

(a) The minimal cut when p is small. There are only isolated voids, so that the cut does not deviate much from vertical and contributes to the entropy like $O(t)$.

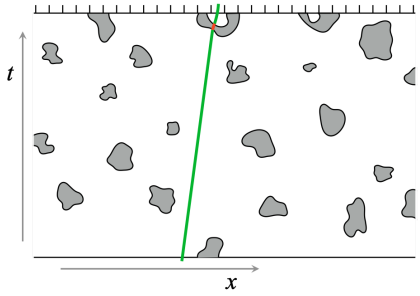

(b) The minimal cut when p is large. The cut can mostly travel through the large connected void, contributing no entropy, except for a small region near the fixed point at the top.

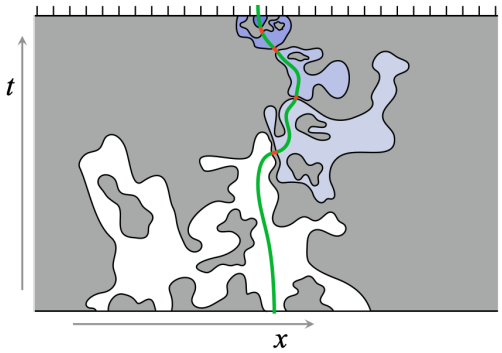

(c) The minimal cut near the critical point $p = p_c$. The cut eventually escapes to a large void after crossing $O(\log t)$ bonds.

FIG. 3: The minimal cuts for the percolation problem at three different values of p. In the large-time limit, the lattice becomes very fine when viewed at the scale t and is thus shown as a continuous plane. The greyed-in areas are connection components of filled-in bonds, which fill most of the plane at small p but become disconnected near $p = p_c$ and eventually become small islands at large p. Similarly, the white regions are voids with no filled-in bonds. Adapted from [1]

III. A SYMMETRY-PROTECTED TOPOLOGICAL DISENTANGLING PHASE

Now we understand that, at least for a simple random circuit in one dimension, entangling and disentangling phases can occur, and we know at least the broad strokes of the entanglement scaling of these states. What other interesting properties do they possess, if any?

In this section as well as the next, we will see that each phase can be made interesting in its own right, and that other phases can also exist under different measurement scenarios. In particular, the disentangling phase can have non-trivial topological properties that

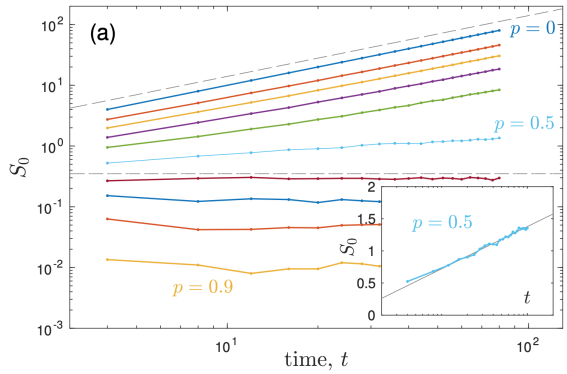

(a) Growth of $S_0(t)$, which is linear for $p < p_c$ and zero for $p > p_c$.

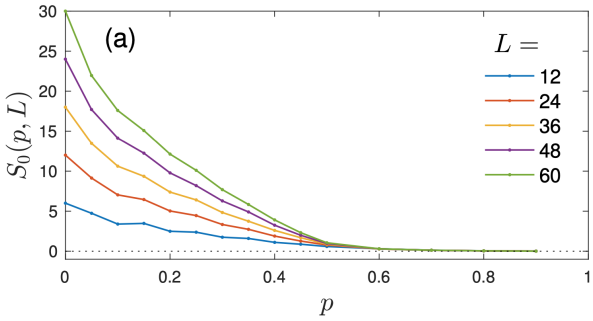

(b) Asymptotic value of S_0, which is macroscopic for $p < p_c$ and near zero for $p > p_c$.

FIG. 4: Numerical simulations of S_0 in the percolation problem for various circuit widths L, showing agreement with the theoretical predictions discussed above. Adapted from [1].

make it distinct from the standard product state, and the entangling phase has properties reminiscent of quantum error correction. We start with the symmetry-protected topological disentangling phase.

Topologically protected states are those that cannot be transformed into one another by only local operators. They are "robust" against the types of noise that are plague real quantum systems, like small random magnetic fields that introduce terms like σ^x on a single site, in the sense that if you act with a local operator (restricted to some allowed set of *stabilizers*) you know that your state will not be destroyed. This makes topologically protected states useful for applications like quantum computing, where if your qubits are in a topologically protected state, they will remain in that state until acted upon with a non-local operator, something that is unlikely to happen naturally. A famous example of topologically protected states is Kitaev's toric code, where the number of protected states depends on a topological feature (the genus) of the underlying manifold. [3]

A. A Symmetry-Protected Model

The reference [4] considered a new model that still relies on block-circuits but has several key difference to the one we've been working with so far. First, it fixes a Hamiltonian H_0

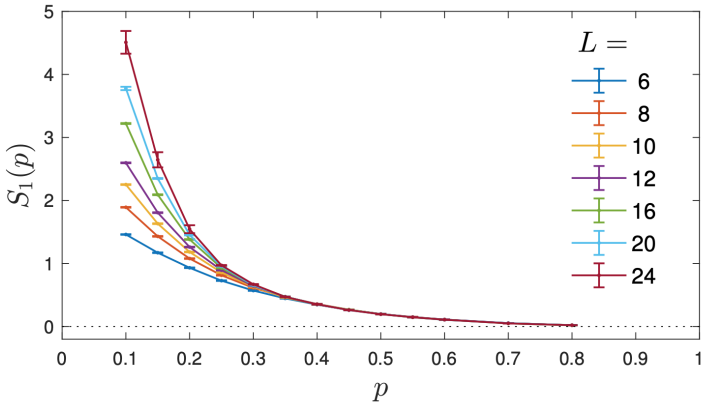

FIG. 5: Asymptotic value of the Von Neumann entropy S_1 as a function of p for various circuit widths L. Note the qualitative similarities to the S_0 case depicted in Figure 4b despite the different value of $p_c \simeq 0.25$. Adapted from [1].

and a $\mathbf{Z}_2 \times \mathbf{Z}_2$ symmetry generated by the operators G_1 and G_2, defined by

$$H_0 = -\sum_{j=2}^{n} \sigma^x_{j-1}\sigma^z_j\sigma^x_{j+1} \tag{III.1}$$

$$G_1 = \prod_{j \text{ even}} \sigma^z_j \qquad G_2 = \prod_{j \text{ odd}} \sigma^z_j \tag{III.2}$$

The eigenstates of these operators, which all commute and so are simultaneously diagonalizable, which defines four new subsets of the Hilbert space indexed by the two eigenvalues of G_1 and G_2 (which are ± 1). These eigenstates are topologically protected by the G's, or simply "symmetry-protected topological" (SPT), because they cannot be transformed into one another by a local perturbation such as a spin flip. Note that many product states of spins are eigenstates of the G's, but only superpositions of these states can be eigenstates of H_0 due to the presence of the σ^x terms. The stabilizers are the terms of H_0, denoted $g_j = \sigma^x_{j-1}\sigma^z_j\sigma^x_{j+1}$.

B. The SPT Phase Transition

To obtain the SPT phase out of a random block circuit, some changes to the structure of the problem must be made. First, we introduce a new type of measurement operator: the stabilizers at each qubit, g_j. Next, we change the schema of the timestep. Instead of a layer of two-qubit unitaries followed by measurement with probability p, at each qubit we either apply a three-qubit unitary (the reason why two-qubit vs. three-qubit is not important) with probability p_u, perform a single-qubit measurement of σ^z with probability p_s, or measure the stabilizer g with probability $1 - p_u - p_s$. A sample circuit following this schema is shown in Figure 6a.

Now there are two tunable parameters, p_u and p_s, and we can map out the phase diagram in the $p_u - p_s$ plane. The result is shown in Figure 6b. Note that the SPT phase actually satisfies an area-law for its entropy scaling, which qualifies it as a disentangling phase in our

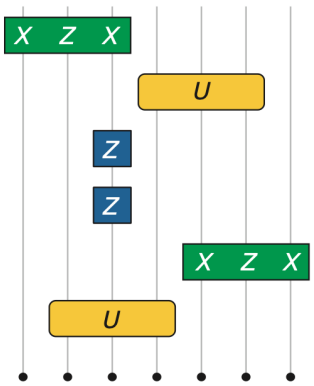

(a) A sample circuit with unitaries, single-site measurements, and stabilizer measurements.

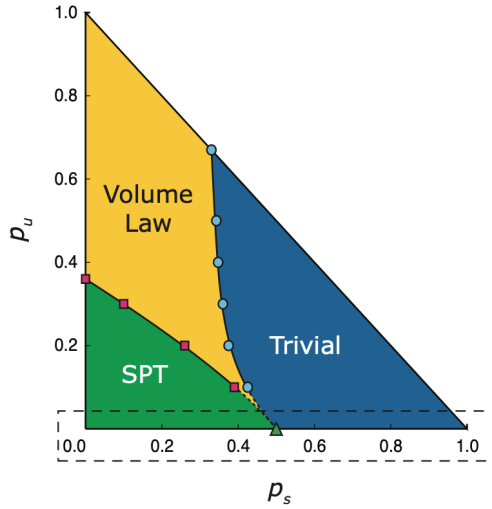

(b) The phase diagram of the SPT system, showing the emergence of an SPT phase.

FIG. 6: Circuit model and phase diagram of the SPT system. The "trivial" phase is a product-state disentangling phase, and while the SPT phase is also disentangling, it is topologically distinct from the trivial phase. Adapted from [4].

taxonomy.

C. The SPT Phase and Error Correction

The emergence of the SPT phase bears some connection to the implementation of active quantum error correction (QEC). Quantum error correction in general seeks to mitigate the effects of noise on the evaluation of quantum circuits by controlling how that noise modifies the state of the system as a function of time. In active QEC, this is accomplished by fixing a set of stabilizers and an associated set of code words, which are just product states. The stabilizers are repeatedly measured over the course of a computation and, while they do not modify the code words, they are able to detect (based on the result of their measurement) if any of the qubits have been flipped by noise. If so, the operator of the computer can then flip the qubit back to actively correct the error detected by the stabilizer.

This is similar to the process of the random formation of the SPT phase, which happens when the rate of stabilizer measurements is high enough compared to the rate of unitary dynamics and single-site measurements. In fact, we can model the unitary dynamics as contributing to continuous-time error accumulation in the code-word represented by the SPT state, and the single-site measurements as faulty applications of stabilizer measurements, as noted in [4]. The presence of a phase transition away from the SPT phase indicates that if the fault rate becomes too high, or likewise with the rate of error accumulation, the quantum code will no longer function because the SPT phase, which we imagine as encoding some sort of logical information in its topological structure, is destroyed.

IV. THE ENTANGLING PHASE, AND ERROR CORRECTION, AND SCRAMBLING

Having seen that the disentangling phase of these circuits can possess some information-protection capability, it is natural to wonder whether the entangling phase could be used for the same purpose. Somewhat unsurprisingly, the answer is yes, and it relates to the defining ability of the entangling phase – guarding against the destruction of quantum information by the measurements, as explicated in [5].

Consider the following asymmetry in the role of the unitaries and projections: given a subsystem A, a unitary can only build entanglement if it acts across the boundary of A, but given a highly entangled state, a projective measurement *anywhere* in the system can destroy entanglement with A. This simple argument seems to suggest that the volume-law phase is not stable under the action of any finite rate of measurements, but as we have seen before, this is empirically not the case. Where is the flaw in the reasoning?

The solution is that the volume-law phase generated by these unitaries has a particular property: it is *information scrambling*, in the sense that the information about the correlations between A and the rest of the system are spread out into highly non-local degrees of freedom that cannot be destroyed by sparse measurements. This is very similar to how some forms of QEC work, and turns out to be a generic feature of quantum dynamics in locally interacting systems.

A. The Decoupling Theorem

To see why the degree of scrambling is important, consider the toy model presented in [5]. The basic situation is shown in Figure 7, where a large number $2N$ of qubits are partitioned into two sets A and B and then individually subjected to unitary evolution by U_A and U_B, which of course do not entangle the subsystems with each other. We control the amount of entanglement initially present between A and B by specifying a fraction γ such that γN of the qubits in A are members of Bell pairs with qubits in B, which are a maximally entangled state. Similar to what we've considered previously, a fraction p of the qubits in A are then measured. The Bell pair qubits carry correlations because of their entanglement, and we will attempt to track the trajectory of these correlations through the system to see if they are destroyed by the projective measurements.

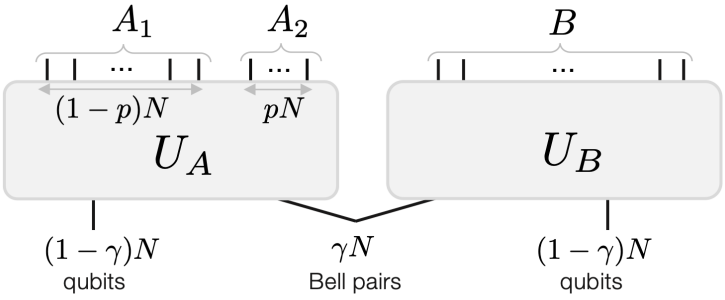

FIG. 7: The setup we use to determine how well the unitaries can scramble the information in the incoming γN Bell pairs. Adapted from [5].

Without loss of generality we can focus on measurement of subsystem A (the results for B will be identical by symmetry). The unitary U_A is randomly chosen, which means that it is generically highly-entangling within the subsystem A. We expect, then, that the correlations with B, initially carried by γN of the qubits, will become quite spread out through all N qubits in A.

To make this precise, let A_1 consist of the $(1-p)N$ qubits that are unmeasured, A_2 consist of the pN qubits that are, and $\widetilde{B}$ consist of the entangled system of A_2 and B. we employ a decoupling theorem [6] that bounds exactly how far the exact density matrix of the composite system $\rho_{A_2\widetilde{B}}$ can get from the "maximal scrambling" case. In this setup, maximal scrambling means the qubits in A_2 are highly entangled with those in $\widetilde{B}$ before the measurement, so that after the measurement, the reduced density matrix ρ_{A_2} of the measured subsystem is in a maximally mixed state, which we denote by $\rho_{A_2}^{\text{mixed}}$. Thus, the maximal scrambling state of $A_2\widetilde{B}$ is

$$\rho_{A_2\widetilde{B}}^{\max} = \rho_{A_2}^{\text{mixed}} \otimes \rho_{\widetilde{B}}$$

Now, assume the measurement rate p is not too large. Then, the decoupling theorem states that, averaged over the random options for U_A, the distance between $\rho_{A_2\widetilde{B}}(U_A)$ and this maximal scrambling state is bounded like so:

$$\mathbf{E}_{U_A}[||\rho_{A_2\widetilde{B}}(U_A) - \rho_{A_2\widetilde{B}}^{\max}||_1] \leq 2^{-(1-2p-\gamma)N/2} \tag{IV.1}$$

where $\mathbf{E}_{U_A}[||\cdot||_1]$ is the expectation of the L_1 norm over the Haar measure that U_A is drawn from. The bottom line is that, in the many-qubit limit of $N \gg 1$, information is scrambled exponentially well if $1-2p-\gamma > 0$, or

$$\gamma < 1 - 2p \tag{IV.2}$$

This inequality is actually not tight, and can be improved to $\gamma < 1-p$ by making use of the classical information obtained from the pN qubit measurements.

Generalizing a bit from the specific setup above, γ can be thought of as an average entropy per incoming qubit at each layer of the block circuit. This scrambling inequality says the entropy density γ should grow to a maximum value of $1-p$ when p is not too large, which means that the total entropy is an extensive quantity under these circumstances. Thus, the volume-law phase emerges.

B. Numerical Evidence for Scrambling in Block Circuits

The authors in [5] also performed numerical experiments to see if, as we expect from the analytic argument above, the rate of scrambling affects the emergence of the volume-law entangling phase. They did this by modifying the block circuits that we have been discussing thus far by replacing the random unitary blocks, which we know to be good scramblers of information, with new blocks that only scramble at a certain rate. How they accomplish this is not particularly important to understand; suffice it to say that a new parameter $d/m \geq 0$ is introduced, with $d/m = 0$ corresponding to no scrambling and $d/m \geq 2$ corresponding to the random unitary scrambling discussed above. By tuning this parameter along with the measurement rate p, the phase diagram in Figure 8 was generated. These results clearly

show that smaller scrambling rates suppress the formation of the volume-law entangling phase.

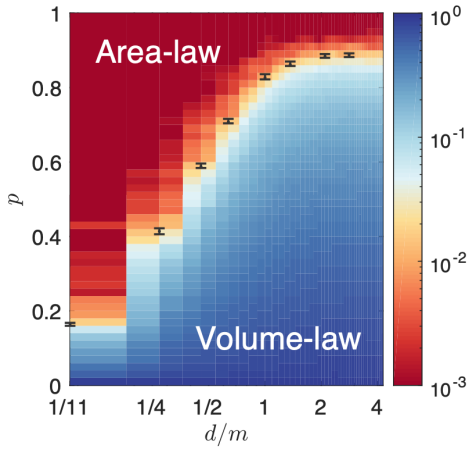

FIG. 8: The $p - d/m$ phase diagram, showing that the area-law phase dominates when scrambling is low and the volume-law phase dominates when scrambling is high. The colorbar is proportional to the steady-state entanglement entropy of the system. Adapted from [5].

V. DISCUSSION AND OUTLOOK

We have seen how certain types of simple, random unitary quantum circuits punctuated by projective measurements can give rise to new phases of quantum matter, where the order parameter is related to the degree of entanglement within the state that the circuit generates. We were given two different perspectives on how these phases emerge: one from a connection to a problem in classical network theory, which indicated that a phase transition ought to occur in the system, and one from the viewpoint of quantum information theory, which emphasized the degree to which information scrambling is important to the stability of the volume-law entangling phase. We also saw that the disentangling phase can be made interesting as well, and connected its emergence to the performance of quantum error correcting stabilizer codes. All of these phenomena were predicted by analytical arguments (with varying degrees of rigor) but confirmed unambiguously by numerical experiments.

What has been missing thus far in the literature? For one, all the systems described thus far have been discrete-time quantum circuits. One interesting question is: to what extent does our entangling/disentangling taxonomy survive when we consider continuous time evolution under a locally interacting model, especially in $D > 1$? This could be an interesting direction for future research, especially given that this sort of local model has a good chance of being implemented directly in real AMO experiments in the near future. Moreover, much effort has been expended on work identifying and characterizing many-body localization (MBL) in these systems, which also exhibits area-law entanglement (among other interesting properties). Very little is currently known about the MBL transition outside of a few model systems, so any connections to the entangling-disentangling transition would likely be warmly welcomed by the condensed matter community.

[1] Brian Skinner, Jonathan Ruhman, and Adam Nahum, "Measurement-Induced Phase Transitions in the Dynamics of Entanglement," Physical Review X **9**, 31009 (2019), arXiv:1808.05953.

[2] J. T. Chayes, L. Chayes, and R. Durrett, "Critical behavior of the two-dimensional first passage time," Journal of Statistical Physics **45**, 933–951 (1986).

[3] A. Yu Kitaev, "Fault-tolerant quantum computation by anyons," Annals of Physics **303**, 2–30 (2003), arXiv:9707021 [quant-ph].

[4] Ali Lavasani, Yahya Alavirad, and Maissam Barkeshli, "Measurement-induced topological entanglement transitions in symmetric random quantum circuits," Nature Physics **17**, 342–347 (2021), arXiv:2004.07243.

[5] Soonwon Choi, Yimu Bao, Xiao Liang Qi, and Ehud Altman, "Quantum Error Correction in Scrambling Dynamics and Measurement-Induced Phase Transition," Physical Review Letters **125**, 30505 (2020), arXiv:1903.05124.

[6] John Preskill, "Chapter 10: Quantum Shannon Theory," PH219/CS219 Lecture Notes **10**, 96 (2016), arXiv:1604.07450.