Random insights into the complexity of two-dimensional tensor network calculations

Sofía González-García [1,2,3,*] Shengqi Sang,[1,2] Timothy H. Hsieh,[1] Sergio Boixo [3] Guifré Vidal,[3]
Andrew C. Potter,[4] and Romain Vasseur[5]

[1]*Perimeter Institute* for Theoretical Physics, Waterloo, Ontario, Canada N2L 2Y5
[2]*University of Waterloo*, Waterloo, Ontario, Canada N2L 3G1
[3]*Google Quantum AI, Santa Barbara, California 93111, USA*
[4]*Department of Physics and Astronomy, and Quantum Matter Institute, University of British Columbia, Vancouver, BC, Canada V6T 1Z1*
[5]*Department of Physics, University of Massachusetts, Amherst, Massachusetts 01003, USA*

Projected entangled pair states (PEPS) offer memory-efficient representations of some quantum many-body states that obey an entanglement area law and are the basis for classical simulations of ground states in two-dimensional (2d) condensed matter systems. However, rigorous results show that exactly computing observables from a 2d PEPS state is generically a computationally hard problem. Yet approximation schemes for computing properties of 2d PEPS are regularly used, and empirically seen to succeed, for a large subclass of ("not too entangled") condensed matter ground states. Adopting the philosophy of random matrix theory, in this work, we analyze the complexity of approximately contracting a 2d random PEPS by exploiting an analytic mapping to an effective replicated statistical mechanics model that permits a controlled analysis at a large bond dimension. Through this statistical-mechanics lens, we argue that (i) although approximately sampling wave-function amplitudes of random PEPS faces a computational-complexity phase transition above a critical bond dimension, and (ii) one can generically efficiently estimate the norm and correlation functions for any finite bond dimension. These results are supported numerically for various bond-dimension regimes. It is an important open question whether the above results for random PEPS apply more generally also to PEPS representing physically relevant ground states.

I. INTRODUCTION

Tensor network states (TNS) provide a compact representation of quantum states with spatially local entanglement, such as ground states of local Hamiltonians. One-dimensional (1d) TNS, matrix product states (MPS), can be efficiently contracted, enabling dramatic progress in studying 1d many-body ground states [1–4]. Projected entangled pair states (PEPS) [5] are a higher dimensional generalization of MPS tensor networks. Contrary to their 1d counterpart, contracting PEPS is a #P-complete task [6] (the complexity class of hard counting problems) even for simple square lattice PEPS with constant bond dimension, D. Moreover, approximately contracting a 2d PEPS with bond dimension $D \sim \text{poly}(L)$ that scales polynomially in the linear dimension L of the system was shown to be average-case hard even for calculating simple physical quantities like expectation values of local observables [7]. Yet, in contrast to these complexity results, the practical experience of PEPS practitioners [8–13] is that the standard algorithm [5,14] (which we review below) for approximating the physical properties of PEPS seems to work efficiently for finitely correlated ground states of 2d lattice models in many condensed matter problems.

Systematically clarifying the behavior of PEPS and their approximate calculation efficiency is crucial not only for computational physics but also for refining the targets for quantum computational advantage in materials and chemistry simulation problems [15], which are one of the leading prospective applications for quantum computers [16]. A key challenge is a lack of systematic analytical tools for analyzing the complexity of approximate contraction schemes. To this end, we adopt the philosophy of statistical mechanics and random matrix theory which show that, in many instances [17], the statistical properties of random ensembles can be systematically understood, even when individual cases cannot be directly analyzed.

With this motivation in mind, we study the "typical" contraction complexity of 2d PEPS with *random* tensors. This enables us to systematically address the statistical complexity of random PEPS through complementary analytical and numerical approaches. Analytically, we exploit a replica-trick-based mapping between entanglement features of random disordered PEPS and partition functions of classical lattice-"spin" models. We present evidence that standard PEPS contraction schemes can efficiently approximate correlations of a 2d random PEPS, including both of local observables and many nonlocal observables such as string order parameters. This contrasts the behavior of computing global properties such as overlaps $\langle \Psi' | \Psi \rangle$ between two PEPS states, which includes individual wave-function amplitudes $\langle s_1, s_2, \dots | \Psi \rangle$, where one expects a
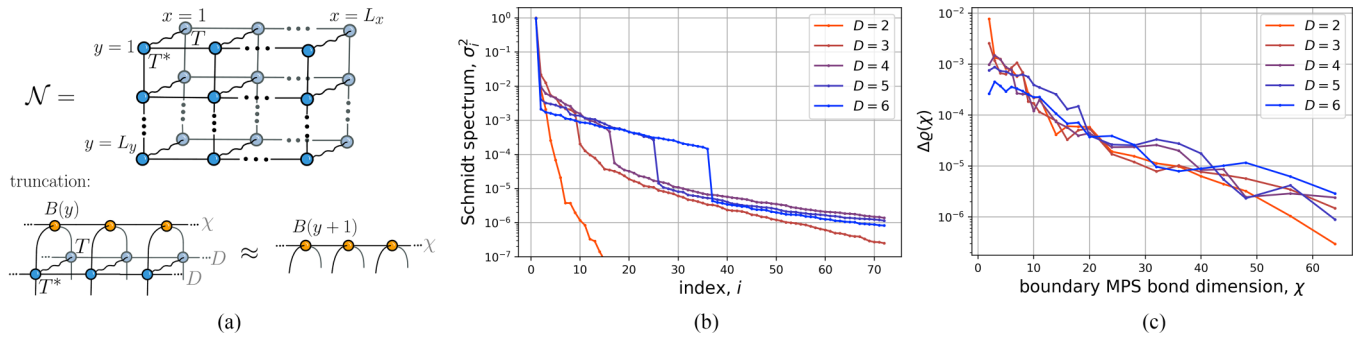
*Corresponding author: sofiagonzalezgarcia@ucsb.edu

FIG. 1. Boundary MPS contraction of 2d PEPS. (a) Tensor network representation of the norm $\mathcal{N} = \langle \Psi | \Psi \rangle$, for a finite $L_x \times L_y$ PEPS $|\Psi\rangle$ (top) and one iteration of the infinite bMPS approximate contraction algorithm (bottom): a row of TT^* tensors is multiplied by the boundary MPS whose bond dimension is then truncated back to χ. (b) Schmidt coefficients at a bipartition of the infinite bMPS for $\chi = 72$ for a clean random iPEPS for bond dimension $D \in \{2, 3, 4, 5, 6\}$. (c) Change $\Delta \varrho(\chi)$ in the physical single-site density matrix $\varrho(\chi)$ as a function of χ, where we take $\varrho(\chi_{\max} = 72)$ as a reference.

complexity-phase-transition tuned by D [18], which has been observed numerically [19,20]. We emphasize that the arguments presented are not rigorous proofs, but rather physical arguments about the replica statistical mechanics that are grounded in standard paradigms of statistical mechanics and critical phenomena.

To validate the predictions of this analytic derivation, we numerically explore the D dependence of contracting both "clean" (translation-invariant) infinite PEPS (iPEPS) [14], as well as "clean" and "disordered" (see below) stabilizer (Clifford) PEPS, which allow for a much larger bond dimension. The numerical results are consistent with the analytic predictions, including quantitative agreement in the detailed large-D asymptotic behavior. While the analytical findings are limited to disordered PEPS, the numerical results suggest that the predictions also hold for "clean" translation-invariant PEPS.

It remains an important open question to determine the relevance of these random PEPS results to PEPS representing ground states of Hamiltonians relevant to condensed matter, materials science, and chemistry. For example, it is known that despite being highly entangled, at a large D, the random square-lattice PEPS exhibit only short-range correlations for local observables [21] (see also Appendix C 5), whereas the ground states of physical systems can exhibit relatively long correlation lengths, potentially making them harder to contract. On the other hand, the statistical mechanics mapping applies more broadly to a large class of network structures, including examples with arbitrarily long-range correlations (see Appendix D 2).

II. PEPS AND BOUNDARY MPS METHOD

We begin by briefly reviewing standard schemes to approximately contract 2d PEPS. Consider a 2d square lattice made of $L_x \times L_y \equiv N$ sites, $L_x \leqslant L_y$, where each site is represented by a d-dimensional vector space $\mathbb{C}^d$, and a *bulk* state $|\Psi\rangle \in \bigotimes_{i=1}^{N} \mathbb{C}^d$ of the lattice that accepts a PEPS representation of the form $|\Psi\rangle = \sum_{s_1,\ldots,s_N=1}^{d} \mathcal{C}[T_{[1]}^{s_1} \ldots T_{[N]}^{s_N}]|s_1 \ldots s_N\rangle$. Here, $T_{[r]}$ denotes the five-index PEPS tensor on lattice site $r \in \{1, 2, \ldots, N\}$, with complex components $(T_{[r]})_{ijkl}^{s}$, where $s = 1, \ldots, d$ and $i, j, k, l = 1, \ldots, D$ are known as *physical*

and *bond* indices, respectively, and $\mathcal{C}$ denotes contraction of the bond indices connecting nearest neighbor tensors.

An archetypal PEPS computation is the contraction of the 2d tensor network in Fig. 1(a), of size $L_x \times L_y$, which corresponds to the overlap $\mathcal{N} \equiv \langle \Psi | \Psi \rangle$. Other closely related 2d networks, such as those corresponding to k-point correlators of local operators $\langle \Psi | O_1 O_2 \ldots O_k | \Psi \rangle$, are contracted similarly. An exact contraction incurs a computational cost $O(L_y \exp(L_x))$, but a number of efficient algorithms exist, most notably the boundary MPS (bMPS) method [5,14], which scales linearly both in L_x and L_y at the price of introducing approximations.[1] The bMPS method considers the *virtual* 1d lattice made of L_x sites, each described by a D^2-dimensional vector space $\mathbb{C}^{D^2}$ (corresponding to two PEPS bond indices), obtained from a horizontal cut of the 2d tensor network $\mathcal{N}$, and a so-called *boundary state* $|\psi\rangle\rangle \in \bigotimes_{i=1}^{L_x} \mathbb{C}^{D^2}$ that accepts an MPS representation, $|\psi\rangle\rangle = \sum_{\alpha_1,\ldots,\alpha_{L_x}} \mathcal{C}(B_{[1]}^{\alpha_1} \cdots B_{[L_x]}^{\alpha_{L_x}})|\alpha_1 \cdots \alpha_{L_x}\rangle\rangle$. Here the double brackets $|\rangle\rangle$ emphasize that the boundary state is actually a vectorized density operator on the PEPS bond space, $B_{[x]}$ is a 3-index MPS tensor on site $x \in \{1, 2, \ldots, L_x\}$ with complex components $(B_{[x]})_{\beta\gamma}^{\alpha}$, where $\alpha = 1, \ldots, D^2$ and $\beta, \gamma = 1, \ldots, \chi$, and χ is the bMPS bond dimension. In order to approximately contract the 2d tensor network $\mathcal{N}$, which we now regard as made of L_y rows of tensors where each row is labeled by an integer $y = 1, \ldots, L_y$, we first represent the boundary state $|\psi(1)\rangle\rangle$ corresponding to the top row ($y = 1$) exactly as a bMPS with bond dimension D^2, which we then truncate down to χ (if $D^2 > \chi$). For increasing value of $y = 1, 2, \ldots$, given a (possibly truncated, approximate) bMPS representation of the boundary state $|\psi(y)\rangle\rangle$, we produce an approximate bMPS for $|\psi(y + 1)\rangle\rangle$ by contracting the row $y + 1$ with the existing bMPS, see Fig. 1(a), then truncating the resulting bond dimension χ' down to χ (if $\chi' > \chi$). Finally, the value of $\mathcal{N}$ is obtained by contracting the bMPS for $|\psi(L_y - 1)\rangle\rangle$ with the bottom row (i.e., $y = L_y$). The leading computational cost of

---

[1]The corner transfer matrix (CTM) [22] method is another prominent approximate contraction scheme for the 2d tensor network $\mathcal{N}$. CTM and bMPS are closely related and similar conclusions regarding boundary entanglement, as discussed in this paper, apply to both.

the bMPS method with respect to index dimensions d, D, χ is $O(D^4\chi^3 + dD^6\chi^2)$ [23] (or $O(D^6\chi^3 + dD^8\chi^2)$ in the infinite case [14], $L_x, L_y \to \infty$).

The above sequential contraction of the 2d tensor network $\mathcal{N}$ is analogous to a dynamical evolution of the state of a 1d system (represented by the bMPS) by a transfer matrix implementing a completely positive map made of a row of TT^* tensors. This setting is similar to 1d noisy quantum circuits [24,25].

III. RANDOM PEPS MODELS

We consider applying the boundary MPS method to random PEPS, whose tensors T are multidimensional arrays of complex entries, with each entry sampled independently and identically distributed (i.i.d.) from the complex normal (Gaussian) distribution with zero mean and unit variance. We consider two different ensembles of random PEPS: (i) "clean," translation-invariant PEPS with a single random instance of T used for every site, and (ii) "disordered" PEPS with a distinct, random $T_{[r]}$ for each site, r. From the perspective of spatial symmetries, clean and disordered random PEPS are akin to crystalline and disordered materials, respectively.

IV. BOUNDARY MPS ENTANGLEMENT AND RANDOM PEPS CONTRACTION

The resulting value of $\mathcal{N} = \langle \Psi | \Psi \rangle$ is approximate due to the errors introduced at each truncation typically implemented by preserving the χ largest Schmidt coefficients $\{\sigma_i\}$ (or coefficients $\{\sigma_i^2\}$ in the entanglement spectrum) assigned to each bond of the bMPS (see Appendix A). Intuitively, if most of the $O(D^{2y})$ Schmidt coefficients (for $y \leqslant L_x/2$) are "large," then an accurate estimate of $\mathcal{N}$ requires using a bMPS with a bond dimension χ that grows exponentially in y, and thus the method is not efficient. However, if only a small number of Schmidt coefficients are significant, then retaining only a constant value of χ may suffice for an accurate bMPS representation, thus enabling efficiently approximate contraction of the network.

Figure 1(b) shows the entanglement spectrum $\{\sigma_i^2\}$ of the *fixed-point* bMPS $|\tilde{\psi}\rangle\rangle$ for clean random iPEPS, where we exploit translation invariance to access the $L_x, L_y \to \infty$ limit, and $|\tilde{\psi}\rangle\rangle$ represents the large-y fixed point of the bMPS sequence $\{|\psi(y)\rangle\rangle\}$ (see Appendix A for further details). Notice the exponential decay of σ_i^2 as a function of i, which implies that an accurate approximation of the boundary state can be obtained with a small bond dimension χ. In turn, this results in an accurate approximation $\varrho(\chi)$ to the reduced density matrix $\varrho = \text{tr}_{\bar{R}}|\Psi\rangle\langle\Psi|/\mathcal{N}$ on a local region R of the 2d lattice and thus also to the expectation value $\langle \Psi|O|\Psi\rangle/\mathcal{N} = \text{tr}(\varrho O) \approx \text{tr}(\varrho(\chi)O)$ for any local observable O supported on R. Indeed, consider the distance $\Delta\varrho(\chi)$ between the density matrix $\varrho(\chi)$ and a reference density matrix $\varrho(\chi_{\max})$ as given by the largest eigenvalue of $|\varrho(\chi) - \varrho(\chi_{\max})|$, where we make the key assumption that $\varrho(\chi_{\max})$ is a good approximation to the exact $\varrho = \varrho(\chi \to \infty)$. As shown in Fig. 1(c) for a single site ϱ, $\Delta\varrho(\chi)$ decays exponentially with χ, with similar behavior observed when R consists of more than one site [see Fig. 12(b)]. We therefore conclude that an accurate approxi-

mation to $\langle \Psi|O|\Psi\rangle/\mathcal{N}$ can be obtained efficiently, for these instances of clean random iPEPS.

In practice, the bMPS entanglement is often characterized using the Renyi entropy of order n,

$$S_A^{(n)} = \frac{1}{1-n}\ln\left(\text{tr}\left[\left(\frac{\rho_A}{\text{tr}\rho}\right)^n\right]\right), \tag{1}$$

where $\rho \equiv |\psi\rangle\rangle\langle\langle\psi|$ and $\rho_A = \text{tr}_{\bar{A}}|\psi\rangle\rangle\langle\langle\psi|$ is the reduced density matrix for a 1d subregion, A, of the evolved bMPS, and the denominator is included to explicitly normalize the state. Specifically, for $n \to 1$, we expect an inefficient contraction when $|\psi(y)\rangle\rangle$ obeys an entanglement *volume law*, where $S_A^{(n)}$ is proportional to the size $|A|$ of region A, and an efficient one when it obeys an entanglement *area law*, where $S_A^{(n)}$ is upper bounded by a constant. An advantage of using a single statistic, such as $S_A^{(n)}$, rather than the full entanglement spectrum is that its ensemble average for $n \geqslant 1$ can be evaluated using a statistical mechanics model, as we will see below.

V. STATISTICAL MECHANICS MODEL MAPPING

The entanglement features of random tensor network contractions can be mapped onto free-energy cost of boundary domains in a classical statistical mechanics (stat-mech) model [18,26–28]. This mapping has been used extensively to study both random holographic tensor networks (tensor networks with physical legs only at the boundary) [18,26], and to explore entanglement growth [29–31], measurement-induced phase transitions [32–38] in random quantum circuits and channels [25]. Here, we adapt it to study the entanglement of the evolved boundary state involved in the contraction of a random, disordered PEPS. In the main text, we merely summarize the key elements of the stat-mech mapping, and refer readers to the supplemental material in Appendix for details. We focus on the calculation of the PEPS norm, $\mathcal{N} = \langle \Psi | \Psi \rangle$, for a disordered, random, square lattice PEPS $|\Psi\rangle$, and will later comment to how this is modified for other observables. As above, we denote the density matrix of the evolved bMPS in the PEPS contraction procedure described above as $\rho = |\psi\rangle\rangle\langle\langle\psi|$.

The mapping exploits a replica trick, $\ln f = \lim_{m\to 0}\frac{1}{m}(f^m - 1)$, to perform averages of nth-Renyi entropies [Eq. (1)]. The objects of interest are averages over powers of the reduced density matrix of the boundary MPS for a region, A: $\mathbb{E}[(\text{tr}\rho_A^n)^m]$, where $\mathbb{E}[\dots]$ denotes averaging over the Gaussian distribution of tensor entries. For integer m, n, this average can be carried out using Wick's theorem. These quantities contain $Q = nm$ copies of ρ. Since each copy of ρ contains $2 \times T$ and $2 \times T^*$ tensors, this results in a sum over all possible Wick pairings of the $2Q = 2mn$ replicas of the tensor network. A given Wick pairing is defined by a set of permutation elements $g_r \in S_{2Q}$ where S_{2Q} is the symmetric group (group of permutations) on $2Q$ elements, for each site r. The resulting entanglement entropy can be written as a free-energy difference [18]:

$$S^{(n)}(A) = \lim_{m\to 0}\frac{1}{m(n-1)}(\mathcal{F}_A - \mathcal{F}_0), \tag{2}$$

where $\mathcal{F}_{A,0} = -\ln \mathcal{Z}_{A,0}$ are the free energies of a classical statistical mechanics model with partition function
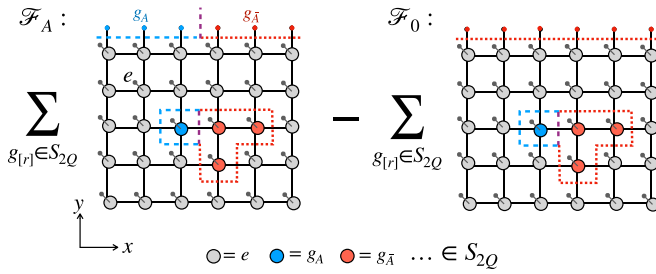
FIG. 2. Statistical mechanics model (schematic). The bMPS entanglement of an random PEPS maps onto the difference between free energies $\mathcal{F}_A - \mathcal{F}_0$ that differ in boundary conditions as shown (note in the replica limit free energy and partition functions coincide). Each site of the random PEPS is replaced by a replica-permutation "spin" (shown as colored circles; while there are $(2Q)!$ different spin values, we show only the three corresponding to the bulk and boundary fields here). Solid lines along bonds of the random PEPS represent ferromagnetic interactions. Lines ending in a dot denote single-site fields. At a large D, the dominant configuration is uniform polarization along e (gray). Fluctuating domains of spins have both a linear boundary tension (dashed lines) due to the interactions, and a surface-area tension due to the bulk e fields.

$\mathcal{Z} = \sum_{\{g_r \in S_{2Q}\}} e^{-H[g]}$, defined by the following Hamiltonian:

$$\mathcal{H}[g] = -\sum_{\langle rr' \rangle} JC(g_r, g_{r'}) - \sum_{r, g \in S_{2Q}} h_{g,r} C(g, g_r). \quad (3)$$

The labels A and 0 in the free energies $\mathcal{F}_{A,0}$ refer to different boundary fields that will be specified below. Here, $\langle rr' \rangle$ denotes neighboring pairs of sites, the function $C(g, g') = C(g^{-1}g')$ counts the number of cycles in the permutation $g^{-1}g'$, and the coupling constant is related to bond dimension: $J = \ln D$. Intuitively, $C(g, g')$ measures how similar the permutations g and g' are to each other, with $C(g, g) = 2Q$, and minimal value $C \geqslant 1$. The coupling between neighboring spins is ferromagnetic, favoring aligned spins. This competes with entropic fluctuations of the spins.

The second term can be thought of as local "fields" that point in a different "directions" (in S_{2Q}) in the bulk and boundary regions. For $\mathcal{F}_A$, the fields are

$$h_{g,r} = \begin{cases} \delta_{g,e} \ln d & r \in \text{bulk} \\ \delta_{g,g_A} \ln D & r \in \text{boundary region } A \,, \\ \delta_{g,g_0} \ln D & r \in \text{boundary region } \bar{A} \end{cases} \quad (4)$$

where e is the identity permutation, and $g_{0,A}$ are permutations corresponding to the trace and cyclic permutation boundary conditions in the entanglement region A and its complement, $\bar{A}$, respectively (see Appendix C for details). In contrast, the boundary fields in $\mathcal{F}_0$ are uniform at the boundary: $h_{g,\text{boundary}} = \delta_{g,g_0} \ln D$ [corresponding to Eq. (4) with only boundary-region $\bar{A}$], see Fig. 2.

VI. GLOBAL OBSERVABLES: COMPLEXITY TRANSITION

For computing global properties of a random PEPS, such as individual wave-function components $\langle s_1 \dots s_N | \Psi \rangle$, or overlaps between distinct random PEPS, $\langle \Psi' | \Psi \rangle$, the tensors are not completely positive, and tensor bulk h_e fields vanish. Absent the bulk $h \sim e$ terms, the model has Q copies of

each tensor T on each site, and a (bulk) left/right $S_Q \times S_Q$ symmetry. There are two possible phases: a disordered phase with intact symmetry, and an ordered phase where this symmetry is spontaneously broken and the spins polarize towards a spontaneously chosen $g \in S_Q$. This model also describes the entanglement of random holographic tensor networks [18], in which physical degrees of freedom arise only at the boundary of a higher-dimensional tensor network with tensors having only virtual bond legs. In the ordered phase, the boundary fields representing the entanglement cut force a g_A to g_0 domain wall at the boundary, which costs energy proportional to the length of the domain, corresponding to volume law scaling of entanglement. By contrast, in the disordered phase, domain walls have a vanishing line tension, leading to area-law entanglement scaling. The universal properties of this complexity transition remain unsolved, but are closely related to those of (forced) measurement-induced entanglement transitions in random unitary circuit dynamics [28]. Numerical analysis [19] for square PEPS indicates that the critical bond dimension for this complexity transition is close to $D_c \approx 2$, meaning that essentially all practical calculations are expected to lie on the hard (volume-law) side of the transitions.

VII. LOCAL OBSERVABLES: AREA-LAW bMPS

For local observables such as the wave-function norm, $\langle \Psi | \Psi \rangle$, or k-point correlators of local operators: $\langle \Psi | O_1 O_2 \dots O_k | \Psi \rangle$, the situation is rather different. Here, the bulk identity-permutation (h_e) fields are present, and explicitly break the left/right $S_{2Q} \times S_{2Q}$ symmetry. This erases the distinction between the ordered and disordered phases of the model, and as we will now argue leaves only the area-law phase behind. In fact, we will see that this conclusion holds even for many nonlocal observables, such as string-order parameters or Wilson loop observables, so long as they lack support on some finite density of sites.

The boundary fields, $g_{A,0}$, associated with the entanglement region A, and its complement region $\bar{A}$, compete with the bulk e fields. The resulting competition is straightforward to analyze at a large D, where domain wall configurations are strongly suppressed, and the partition functions are dominated by the lowest energy configuration. We refer to this approximation as "min-cut." When the L_y is large (to approach the thermodynamic limit, we are interested in $L_y \to \infty$), the number of bulk fields (which scales with the system volume) overwhelms that of the boundary fields (which scale with the linear dimension), forcing the spins to polarize along e. Consequently the linear free energy per unit length of the boundary regions is proportional to $C(g_{A,0}, e) \ln D$ in the $A, \bar{A}$ regions, respectively. Crucially, while the bulk fields explicitly break the left/right $S_{2Q} \times S_{2Q}$, there is a residual bulk $S_Q \times S_Q \subset S_{2Q}$ symmetry that guarantees that $C(g_A, e) = C(g_0, e)$, i.e., that the leading contribution to the boundary-field energies cancels in $\mathcal{F}_A - \mathcal{F}_0$, leading to area-law scaling.

In fact, in the $D \to \infty$ limit, this cancellation is exact, and there is strictly zero operator entanglement. Generically, corrections to the min-cut approximation for S_A come from domains that span across the entanglement cut between A and $\bar{A}$. Since these fluctuations are massive, (the probability of getting a large domain of size ℓ is expected to be

exponentially small in ℓ), these fluctuations are only sensitive to the local change in boundary fields between A and $\bar{A}$, and hence can give only an area-law contribution to S_A. In Appendix C 4, we estimate that, beyond the infinite-D limit, the leading corrections to the area-law entanglement come at order $1/D^2$.

We remark that these results highlight a counterintuitive feature of random tensor networks. While highly correlated states require large bond dimension, *typical* large-D PEPS have rather short-range correlations (despite being quite entangled), and the set of highly correlated large-D PEPS is in fact rare (by the measure of our Gaussian-random ensemble). See Appendix D for a more detailed discussion of potential differences between random PEPS and physical groundstates relevant to condensed matter. In Appendix D 2, we also discuss different tensor network geometries for which *typical* states exhibit longer range correlations.

We emphasize that, while the calculations described were done at leading order in the large-D limit, the general symmetry principles described above suggest that the prediction of area-law entanglement for the boundary-MPS hold also at any D and that there is no expected singular change (phase transition) in the entanglement as a function of D.

VIII. CORRELATORS AND OVERLAPS

While we have focused on the computation of the PEPS norm, this is closely related to the calculation of correlation functions of local operators: $\langle \Psi | O_1 O_2 \dots O_k | \Psi \rangle$. The insertion of local operators amounts to replacing the tensors on those sites with nonpositive tensors, which in the statistical mechanics model corresponds to removing the local e field on that site. Yet, so long as a finite density of sites are unaffected by the operator insertions, there remains a net bulk e field that explicitly breaks the replica-permutation symmetry and removes the complexity phase transition. Hence, we expect our "easiness" result to apply to a very large class of observables including even exotic nonlocal string correlation functions to used detect topological phases and confinement in gauge theories.

IX. ENTANGLEMENT BARRIER AND RANDOM STABILIZER PEPS

We have so far considered the thermodynamic limit of the boundary-MPS entanglement, where the statistical mechanics model predicts that the evolved boundary MPS eventually reaches an area-law steady state. We can also use the stat-mech mapping to examine the transient dynamics of the boundary MPS at early stages in contraction of a PEPS with open boundaries at $y = 0$. For finite y, there is a competition between the bulk fields that favor the spins to point along e, with strength $\sim \ln d$ which contribute an energy that scales with the area of the contracted region (proportional to y), and the boundary fields that favor pointing along g_A or g_0 with strength $\ln D$, but which only arise at the boundary and do not scale with y. When $D \gg d$, comparing the bulk $\sim yL \ln d$ and boundary $\sim L \ln D$ energy costs, we can identify a crossover scale: $\ell^* = \ln D / \ln d$ where these two effects are competitive.

When y is small, $y \ll \ell^* = \ln D / \ln d$, for $D > d$, the boundary fields dominate and the spins g_r polarize along the
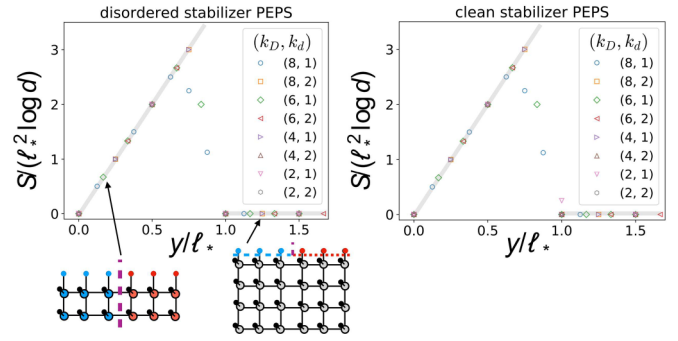

FIG. 3. Entanglement barrier. Mid-cut entanglement entropy $S(y)$ of the bMPS $|\psi(y)\rangle\rangle$ for disordered (left) and clean (right) random stabilizer PEPS as a function of layer number y, for various choices of $(D, d) = (173^{k_D}, 173^{k_d})$. A schematic representation of the dominant spin configurations at a large D in the corresponding stat-mech model are shown underneath the plot. $S(y)$ initially increases linearly up to length scale $\ell_* \sim \ln D / \ln d$ (corresponding to a growing vertical domain wall) before dropping to a small constant for $y \gtrsim \ell_*$ (corresponding to the bulk fields expelling the domain wall out to the system's boundary). In both cases, the simulated $S(y)$ agrees well with the stat-mech model in the large-D regime (solid gray lines), with very small instance to instance variance.

direction of the fields at their closest boundary. For $\mathcal{F}_A$, this adds an extra g_A-to-g_0 domain wall along the y direction with free energy cost $\sim y$ (see Fig. 3). In the limit $D \to \infty$, the bMPS entropy increases linearly with y up to a maximum of $\ln^2 D / d$ at $y = \ell^*$, beyond which, there is a discontinuous jump to 0 bMPS entanglement for all $y > \ell^*$. For a finite D, fluctuations about the minimal domain wall configuration will smooth this jump into a continuous crossover. For a large D, there will still be a pronounced peak in the bMPS entropy at location $y \approx \ell^*$ and with height $\approx \ln^2 D / \ln d$, but the bMPS entanglement beyond the peak saturates to a nonzero value that we estimate as $\mathcal{O}(1/D^2)$ for large-D.

To summarize, as a result of this bulk-boundary crossover, the statistical mechanics model predicts that the entanglement grows as

$$S^A(y) \sim \begin{cases} y \ln D^2 & y \ll \ell^* \\ \text{constant} & y \gg \ell^* \end{cases}, \quad \ell^* \equiv \ln D / \ln d, \quad (5)$$

achieving a maximum of $S_{\max} \sim \ln^2 D / \ln d$ at $y \approx \ell^* = \ln D / \ln d$, corresponding to a maximum bMPS bond dimension of $\chi_{\max} \sim \exp(\ln^2 D / \ln d)$. We emphasize that this calculated scaling form for $\chi_{\max}$ is exact for $D \to \infty$, and corrections from fluctuations can be systematically captured by a low-temperature expansion of the stat mech model when $D \gg d$, and do not change the leading scaling behavior of $\chi_{\max}$ with D, d in that limit.

We note that a similar entanglement barrier arises in the classical simulation of random quantum circuits in the presence of decoherence [24,25], which are relevant to "quantum supremacy" experiments [39]. For circuits, the entanglement barrier reflects the initial build up of quantum correlations and entanglement before noise and errors overwhelm the system making it essentially indistinguishable from a maximally mixed state. The evolution of the boundary MPS is similar to the dynamics of circuits with decoherence (even though

the bMPS evolution is not generally trace-preserving), which provides a physical picture for the low (area law) entanglement of the MPS steady state [24,25]. In this correspondence, the PEPS physical dimension d represents the strength of noise, and the barrier height decreases with increasing noise strength.

We remark that, in practice, the entanglement barrier may not be an obstacle for extracting bulk properties of a PEPS. As illustrated by the iPEPS numerics above, it may not be necessary to accurately track the transient evolution of the boundary MPS through the entanglement barrier. Instead, one may be able to directly find an accurate representation of the area-law steady state bMPS, with χ reflecting the much smaller steady-state entropy deep in the bulk. However, for select tasks, such as simulating edge states of topological materials, the entanglement barrier may have practical consequences. We also note that this prediction is consistent with the average-case hardness proof for computing norms of finite-size random PEPS [7] with $D \sim \text{poly}(L)$: in the stat-mech description this corresponds to an entanglement barrier with $\chi_{\max}$ scaling in a (very weakly) superpolynomial fashion with L. We emphasize that the results of Ref. [7] showing that the average case is #P-complete assumes essentially *exact* contraction with errors that are at most exponentially small in system size. By contrast, our large-D calculations above are relevant for constant (system-size independent) error ϵ. In practice, constant error is a much more reasonable requirement for typical condensed matter calculations—local correlation functions cannot be sensitive to global orthogonality catastrophe effects, a fact that is born out in practical experience. While we cannot, at present, rigorously bound the errors in observables, we expect that demanding error $\epsilon \sim e^{-L}$ would require increasing this $\ln \chi_{\max} \sim \ln^2 L$ result by a polynomial factor of L, which would result in exponential-in-L contraction complexity.

X. STABILIZER PEPS

We can also use the entanglement barrier prediction simply as a test for the validity of the min-cut approximation and replica limit for the statistical mechanics model. To easily access a large D, we focus on random stabilizer (Clifford) PEPS, which can be efficiently simulated via the stabilizer formalism with $\sim \text{poly}(\ln D)$ complexity (see the Appendix B for the definition of stabilizer PEPS and the algorithm for contracting them). For stabilizer PEPS, each of T_i's bond dimensions needs to be an integer power of some prime number p: $d = p^{k_d}$ and $D = p^{k_D}$. In simulations, we reach the large-D regime by picking $p = 173$ (and confirm that qualitatively similar results also hold for $p = 2$).

We performed simulations of both disordered and clean stabilizer PEPS with periodic boundary conditions, and show results in Fig. 3. The simulation results show the entanglement barrier behavior with the characteristic cross-over scale ℓ_* predicted by Eq. (5). Namely, the curves for $S/\ell_* \ln d$ as a function of y/ℓ_* for different d, D collapse (without any adjustable parameters) onto a single universal scaling form that agrees quantitatively with (5), except in a narrow crossover region near $y \approx \ell_*$. Upon increasing D, this crossover sharpens towards the discontinuous jump predicted by (5) for the $D = \infty$ limit. This suggests the validity of the statistical mechanics

model in capturing the entanglement features of the boundary state. The agreement for the clean case is noteworthy since the statistical mechanics model was derived for disordered PEPS, and suggests that the conclusions of the stat-mech model also apply to translation invariant PEPS.

XI. DISCUSSION

We have presented numerical and analytical evidence that the complexity of (approximate) 2d random PEPS calculations greatly depends on the nature of the object to be contracted. On the one hand, overlaps between a 2d random PEPS and a product state $\langle s_1 \cdots s_N | \Psi \rangle$ or, more generally, overlaps $\langle \Psi' | \Psi \rangle$ between two distinct 2d random PEPS, face a complexity phase transition as a function of bond dimension. On the other hand, the overlap $\langle \Psi | \Psi \rangle$ of a single random PEPS as well as (computationally related) physically relevant properties such as the expectation value $\langle \Psi | O | \Psi \rangle$ or correlation functions $\langle \Psi | O_1 O_2 \cdots O_k | \Psi \rangle$ of local operators, can be efficiently computed. The statistical mechanics model provides an intuitive picture for this result.

The chief objective hazard of our analytic analysis is the reliance on the replica trick: away from the infinite-D limit, we are able to analyze properties only for positive integer number of replicas $n = 2, 3, \ldots$ By contrast (see Appendix A), errors in observables are best captured by the von-Neumann entanglement entropy (which can be obtained by taking the limit of Renyi index $n \to 1^+$). We note that stabilizer PEPS have a fine-tuned, completely flat entanglement spectrum and do not constitute an independent check of the replica limit convergence. Therefore, in Appendix A, we numerically explore finite-χ scaling of the n dependence of bipartite Renyi entropies for clean random iPEPS. We observe a smooth and well-converged Renyi-index dependence up to the largest accessible $D = 6$, implying no signs of trouble for the replica limit.

Our statistical mechanics mapping, when generalized to a 3d random PEPS $|\Phi_{3d}\rangle$, predicts again an area-law entangled boundary state for the overlap $\langle \Phi_{3d} | \Phi_{3d} \rangle$ and expectation value and correlation functions of local observables. This 2d area-law entangled boundary state might again be efficiently approximated with a 2d PEPS $|\Psi_{2d}\rangle$. It is thus plausible that approximate expectation values and correlators of local observables can be efficiently evaluated also for 3d random PEPS.

An important question is to what extent the present results obtained for random PEPS may also apply to the contraction of PEPS that represent ground states of interest in condensed matter, materials science, or quantum chemistry. Elucidating such a question would be useful both for the classical simulation of such systems, as well as to determine whether quantum computers may provide an exponential speed-up for such problems [15]. We note that, even if classical TNS methods could efficiently calculate properties of finitely correlated states in 2d and 3d, many challenging tasks remain ripe for quantum advantage, including simulating highly entangled states such as metals or phase transitions, or calculating dynamical responses such as transport coefficients or optical spectra. Progress in addressing this question will likely require
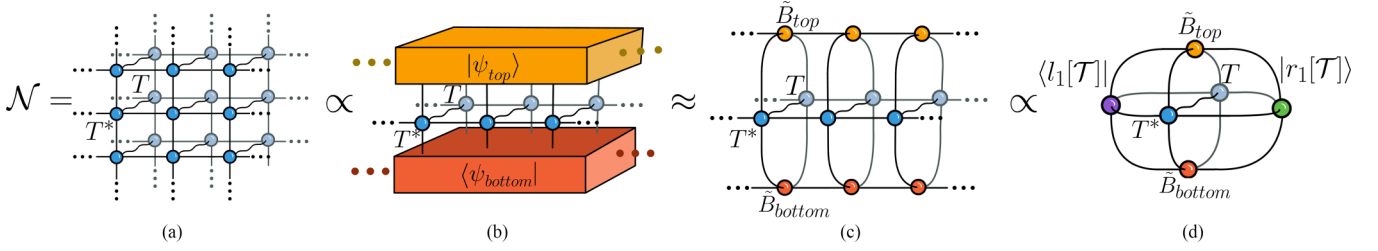
FIG. 4. iPEPS norm computation. (a) Infinite square lattice tensor network of the norm $\mathcal{N}$ of an iPEPS. (b) Norm boundary states $|\psi_{\text{top}}\rangle$ and $\langle\psi_{\text{bottom}}|$ with a 1d transfer matrix in between, which is proportional to the iPEPS norm. (c) The bMPS algorithm obtains an approximation to $|\psi_{\text{top}}\rangle$ and $\langle\psi_{\text{bottom}}|$ with an infinite bMPS. (d) Contracting from the left and from the right of the infinite chain in (c), we obtain the left and right boundary states $\langle l_1[\mathcal{T}]|$, $|r_1[\mathcal{T}]\rangle$ (see later for details on their calculation).

a combination of analytical work and numerical simulations directly targeting "physically relevant" systems of interest.

APPENDIX A: NUMERICAL METHODS: IPEPS

Consider a 2d square lattice made of $L_x \times L_y \equiv N$ sites, $L_x \leqslant L_y$, where each site is represented by a d-dimensional vector space $\mathbb{C}^d$, and a state $|\Psi\rangle$ on the lattice that can be expressed as a PEPS as introduced in the main text. The calculation of the norm $\mathcal{N} = \langle\Psi|\Psi\rangle$ is a common computation in PEPS algorithms, closely related to the computation of more complicated quantities of interest, e.g., expectation values $\langle\Psi|O|\Psi\rangle$ or correlation functions $\langle\Psi|O_1 O_2 \cdots O_k|\Psi\rangle$ of local operators. The exact contraction of the norm of $\mathcal{N}$ has a $O(L_y \exp(L_x))$ computational cost. The boundary MPS method is one of many related approaches [14,22,23,40–42] for approximately computing $\mathcal{N}$, while keeping the computational cost to scale linearly in L_x and L_y. In this Appendix, we consider an infinite PEPS (iPEPS), with $L_x, L_y \to \infty$, by leveraging the translational invariance of the state, and where the computational cost is independent of L_x and L_y.

The tensor network for $\mathcal{N}$ consists of an infinite square lattice with a pair of tensors TT^* on each site, where T and T^* are connected over their physical index [see Fig. 4(a)]. One can show that the norm $\mathcal{N}$ diverges/vanishes with L_x, L_y as $\mathcal{N} = \nu^{L_x L_y}$, where ν is a non-negative real number known as the *norm per site*. Although it is possible to compute ν, in practice this is not needed if we are interested in normalized quantities such as $\langle\Psi|O|\Psi\rangle / \mathcal{N}$ or the reduced density matrix ϱ (Fig. 5), given by the ratio of two 2d tensor networks that only differ by the insertion of (one or more) local operators (see Fig. 6). In this case, we need to compute the fixed point boundary states $|\psi_{\text{top}}\rangle\rangle$, $\langle\langle\psi_{\text{bottom}}|$ (assumed to be unique) of the 1d transfer matrix given by an infinite row of TT^* [see Fig. 4(b)]. From now on we drop the double braket notation $|\rangle\rangle$, $\langle\langle|$ for boundary states, in favor of single braket notation $|\rangle$, $\langle|$. The key idea behind the boundary MPS (bMPS) algorithm is to *approximately* represent the infinite top and bottom boundary states $|\psi_{\text{top}}\rangle$ and $\langle\psi_{\text{bottom}}|$ with two infinite 1d MPS of bond dimension χ. The algorithm detailed below describes how to obtain the top fixed point boundary MPS (the generalization to the bottom boundary is straightforward).
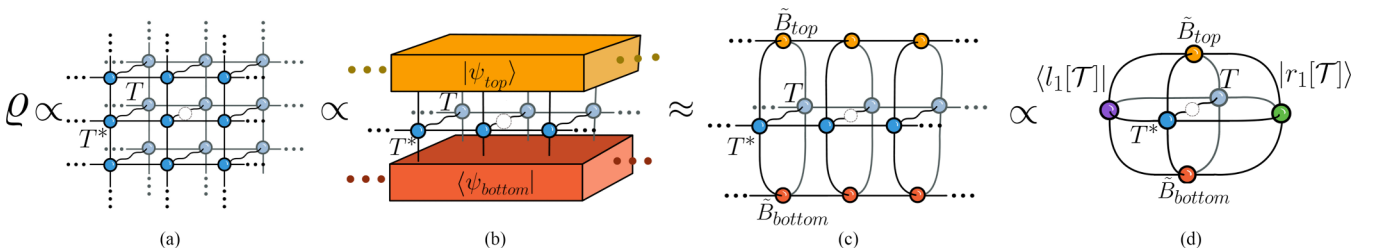

FIG. 5. iPEPS reduced density matrix computation. Computation of the iPEPS 2d reduced density matrix for one site ϱ. It is closely related to $\mathcal{N}$ in Fig. 4 but a physical index is left open. Note that to obtain the correct $\varrho(\chi)$ we need to normalize the object in (d) as in Fig. 6(a).
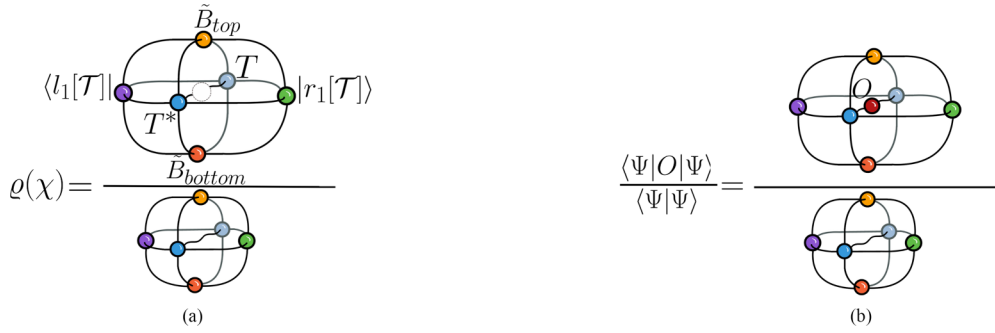
FIG. 6. Computation of observable and reduced density matrix. (a) Tensor netwok for the computation of normalized $\varrho(\chi)$. More detailed calculation of these objects is described below. (b) Tensor network for the computation of observables $\langle \Psi | O | \Psi \rangle / \langle \Psi | \Psi \rangle$.

1. Boundary MPS: method

We build an initial boundary MPS by multiplying pairs of tensors TT^* after removing the vertical upper index (by fixing their value to 1).[2] This translation-invariant bMPS is characterized by a tensor B with 4 indices represented graphically as "legs": 2 bond legs connecting the tensors with each other, and other 2 legs, one joining B with T and the other with T^*. This initial bMPS has bond dimension D^2. If $D^2 \leqslant \chi$, we act with the transfer matrix n times until the first truncation is required (namely, when $D^{2n} > \chi$). After this point, we iteratively perform two operations: (i) truncation of the boundary MPS bond dimension down to χ and (ii) application of the 1d TT^* transfer matrix which increases the bMPS bond dimension to χD^2. This algorithm ends once the bMPS has *converged* (see Appendix A 1 b). The converged bMPS $|\tilde{\psi}\rangle$, characterized by tensor $\tilde{B}$, is our approximation to the true fixed point boundary state $|\psi_{\text{top}}\rangle$.

a. Infinite MPS bond dimension truncation

A key component of the algorithm described above is to perform a bond dimension truncation of the infinite bMPS from $\chi' \to \chi$, where $\chi' > \chi$, while trying to keep a high fidelity between the states represented by $B(\chi)$ and the original $B(\chi')$. In order to do this, we consider a bond bipartition into left and right halves of the infinite spin chain represented by the bMPS. We bring this bond into the canonical form [43], such that the bond index, i, corresponds to the labeling of Schmidt vectors in the Schmidt decomposition of the bMPS representation of the state $|\psi\rangle$ across that index. More explicitly,

$$|\psi\rangle = \sum_{i=1}^{\chi'} \sigma_i \left| \Phi_i^L \right\rangle \otimes \left| \Phi_i^R \right\rangle, \tag{A1}$$

where σ_i label the Schmidt coefficients (normalized such that $\sum_i \sigma_i^2 = 1$ and arranged in decreasing value $\sigma_1 \geqslant \sigma_2 \geqslant \cdots \geqslant \sigma_{\chi'} \geqslant 0$) and $|\Phi_i^L\rangle, |\Phi_i^R\rangle$ are the Schmidt vectors, which form an orthonormal set $\langle \Phi_i^L | \Phi_j^L \rangle = \langle \Phi_i^R | \Phi_j^R \rangle = \delta_{ij}$. The basis $|\Phi_i^L\rangle, |\Phi_i^R\rangle$ are also the eigenvectors of the reduced density

matrices ρ^A and ρ^B, with eigenvalues σ_i^2. Once we perform a gauge transformation to bring the bond into this basis, we can truncate the bond dimension of the bMPS by keeping only the largest χ Schmidt coefficients. We apply this truncation for every bond in the infinite chain. This truncation procedure will keep truncation errors low as long as the entanglement spectrum eigenvalues $\{\sigma_i^2\}$ decay sufficiently fast with i [43,44].

We summarize the truncation algorithm in steps i to vi as depicted in Fig. 7(a). We need to obtain the invertible matrices M and P that bring the translationally invariant bMPS tensor B into the left and right canonical form B_l and B_r, respectively. These gauge requirements are such that the left and right dominant eigenvectors of $B_l B_l^*$ and $B_r B_r^*$ respectively (contracted over the physical bMPS index as shown in Fig. 7(b) should be the identity and its corresponding dominant eigenvalue, λ_1, should be 1 for normalization. These conditions are depicted in Fig. 7(b). Starting from the large bond dimension infinite bMPS that we wish to truncate [Fig 7(a)(i)], we insert two resolutions of the identity $M^{-1}M$ and PP^{-1} at every bond [Fig 7(a)(ii)]. This yields an infinite chain of $P^{-1}BM^{-1}$ tensors denoted B_c [Fig 7(a)(iii)] with a matrix MP on the bonds [Fig 7(a)(iv)]. We consider the singular value decomposition of the matrix $MP = USV^\dagger$. The diagonal entries of the matrix S and the isometries $U, V^\dagger$ correspond to the Schmidt decomposition presented in Eq. (A1) if we had considered only one bond partitioning the infinite chain in half. We truncate these objects and keep only the largest χ singular values and corresponding vectors therefore approximating the bond with smaller sized matrices $MP \approx U'S'V'^\dagger$ [Fig 7(a)(v)]. Putting everything together, as in Fig. 7(a)(vi), we obtain a final expression for $B(\chi) = S'^{1/2}V'^\dagger P^{-1}B(\chi')M^{-1}U'S'^{1/2}$.

b. Convergence of the boundary MPS to a fixed point

After a sufficient number of iterations y_* of the 1d TT^* transfer matrix followed by the truncation as described in Appendix A 1 a, we observe that our boundary MPS has *converged*, meaning that $B(y_*) \approx B(y_* + 1)$ within a set tolerance, in which case we define $\tilde{B} = B(y_* + 1)$. This equivalence can be assessed by looking at the fidelity per site (see Appendix A 2 b for definition) between $B(y)$ and $B(y + 1)$ (which should be approximately 1 if convergence has been reached) or by comparing the Schmidt spectrum of the

---

[2]Because we are interested in the fixed point (assumed to be unique) of the 1d transfer matrix, the choice of initial $|\psi(y = 0)\rangle$ is not important.
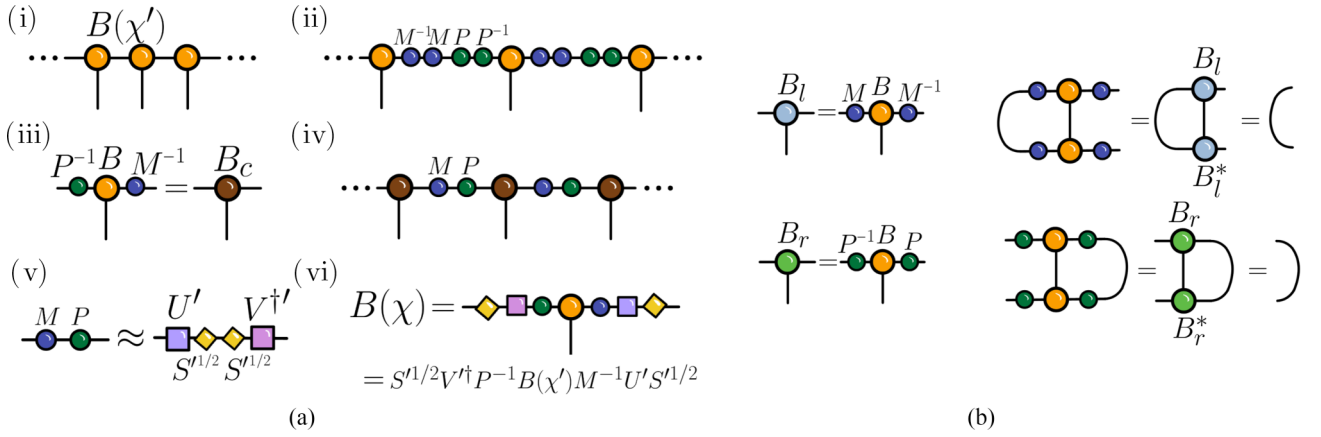
FIG. 7. Truncation of bMPS bond dimension: (a) Steps i–iv in the algorithm for truncating the bond dimension of the bMPS. (b) Finding left and right canonical form gauge transformation matrices M (top) and P (bottom).

boundaries in successive iterations (which should be equal). In our work, we use both metrics to ensure convergence.

2. Infinite MPS and infinite PEPS: computation of metrics

a. Infinite MPS: norm and correlation length

Consider an infinite MPS $|\psi\rangle$ given by tensor B of bond dimension χ. Its transfer matrix $\mathcal{T}_{\langle\psi|\psi\rangle}$ [see Fig. 8(a)] has eigenvalue decomposition:

$$\mathcal{T}_{\langle\psi|\psi\rangle} = \sum_{i=1}^{\chi^2} \lambda_i |r_i\rangle \langle l_i|, \tag{A2}$$

where $\{\lambda_i\}$ is the set of eigenvalues, organized such that $|\lambda_1| > |\lambda_2| \geqslant \cdots \geqslant |\lambda_{\chi^2}|$, and $|r_i\rangle$, $\langle l_i|$ are the right and left eigenvectors, respectively. One can show that: (i) each eigenvalue λ is either real or part of a complex conjugate pair λ, λ^* and (ii) the dominant eigenvalue λ_1 is real and non-negative. We have also assumed that the dominant eigenvalue λ_1 is not degenerate. The norm $\langle\psi|\psi\rangle$ [Fig. 8(a)] is given by $\lambda_1^{L_x}$, where L_x is the number of tensors in the MPS, and thus $L_x \to \infty$ for an infinite MPS. If $\lambda_1 \neq 1$, implying that $\langle\psi|\psi\rangle \neq 1$, we may obtain a normalized MPS $|\psi'\rangle$ with tensor $B' = B/\sqrt{\lambda_1}$ so that $\lambda_1' = 1$ and $\langle\psi'|\psi'\rangle = 1$.

The correlation length ξ is calculated by considering the dominant and second largest eigenvalue λ_1, λ_2 and it is given by

$$\xi = \frac{-1}{\ln(|\lambda_2[\mathcal{T}_{\langle\psi|\psi\rangle}]/\lambda_1[\mathcal{T}_{\langle\psi|\psi\rangle}]|)}, \tag{A3}$$

where we introduce the notation $\lambda_i[\mathcal{T}_{\langle\psi|\psi\rangle}]$ to refer to the ith eigenvalue of the transfer matrix $\mathcal{T}_{\langle\psi|\psi\rangle}$.

b. Infinite MPS: fidelity

Consider two iMPS $|\psi_1\rangle$, $|\psi_2\rangle$ given by tensors B_1, B_2 and bond dimensions χ_1, χ_2 respectively, normalized such that $\langle\psi_1|\psi_1\rangle = \langle\psi_2|\psi_2\rangle = 1$, that is $\lambda_1[\mathcal{T}_{\langle\psi_1|\psi_1\rangle}] = \lambda_1[\mathcal{T}_{\langle\psi_2|\psi_2\rangle}] = 1$ (see previous section). Their fidelity, $F[\psi_1, \psi_2] = \langle\psi_1|\psi_2\rangle$ can be expressed as $F = F_s^{L_x}$, where the *fidelity per site* F_s is such that $0 \leqslant |F_s| \leqslant 1$. One can show that $|F_s| = 1$ if and only if B_2 equals B_1 up to a so-called *gauge transformation*, in which case both tensors give rise to the same state. To calculate F_s consider the transfer matrix given by $B_1 B_2^*$, as depicted in Fig. 8(b). This transfer matrix can be decomposed in the standard way:

$$\mathcal{T}_{\langle\psi_1|\psi_2\rangle} = \sum_{i=1}^{\chi_1 \chi_2} \lambda_i |r_i\rangle \langle l_i|, \tag{A4}$$

where $\{\lambda_i\}$ is the set of eigenvalues, organized such that $|\lambda_1| > |\lambda_2| \geqslant \cdots \geqslant |\lambda_{\chi_1 \chi_2}|$, and $|r_i\rangle$, $\langle l_i|$ the right and left eigenvectors, respectively. Here, we have assumed that the dominant eigenvalue λ_1 is not degenerate. Then, the fidelity per site F_s is simply given by this dominant eigenvalue, $F_s = \lambda_1[\mathcal{T}_{\langle\psi_1|\psi_2\rangle}]$.

c. Infinite PEPS: reduced density matrix

The reduced density matrix for a local region R of the 2d lattice is denoted $\varrho_R = \text{tr}_{\bar{R}} |\Psi\rangle\langle\Psi| / \mathcal{N}$, where the trace is over all sites except those in R. For simplicity, we consider the case where R consists of a single site and refer to this
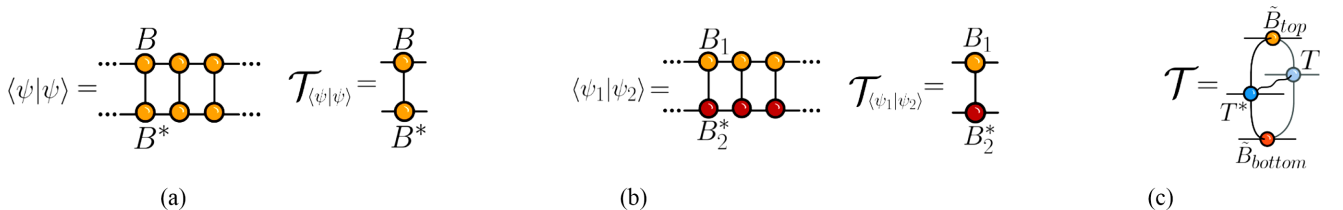

FIG. 8. Transfer matrices. (a) Norm $\langle\psi|\psi\rangle$ of an infinite MPS (left) and the transfer matrix $\mathcal{T}_{\langle\psi|\psi\rangle}$ involved in its calculation. (b) Fidelity $F = \langle\psi_1|\psi_2\rangle$ between two infinite MPS and the transfer matrix $\mathcal{T}_{\langle\psi_1|\psi_2\rangle}$. (c) Transfer matrix $\mathcal{T}$ used for computing $\mathcal{N}$, ϱ and $\langle\psi|O|\psi\rangle$.
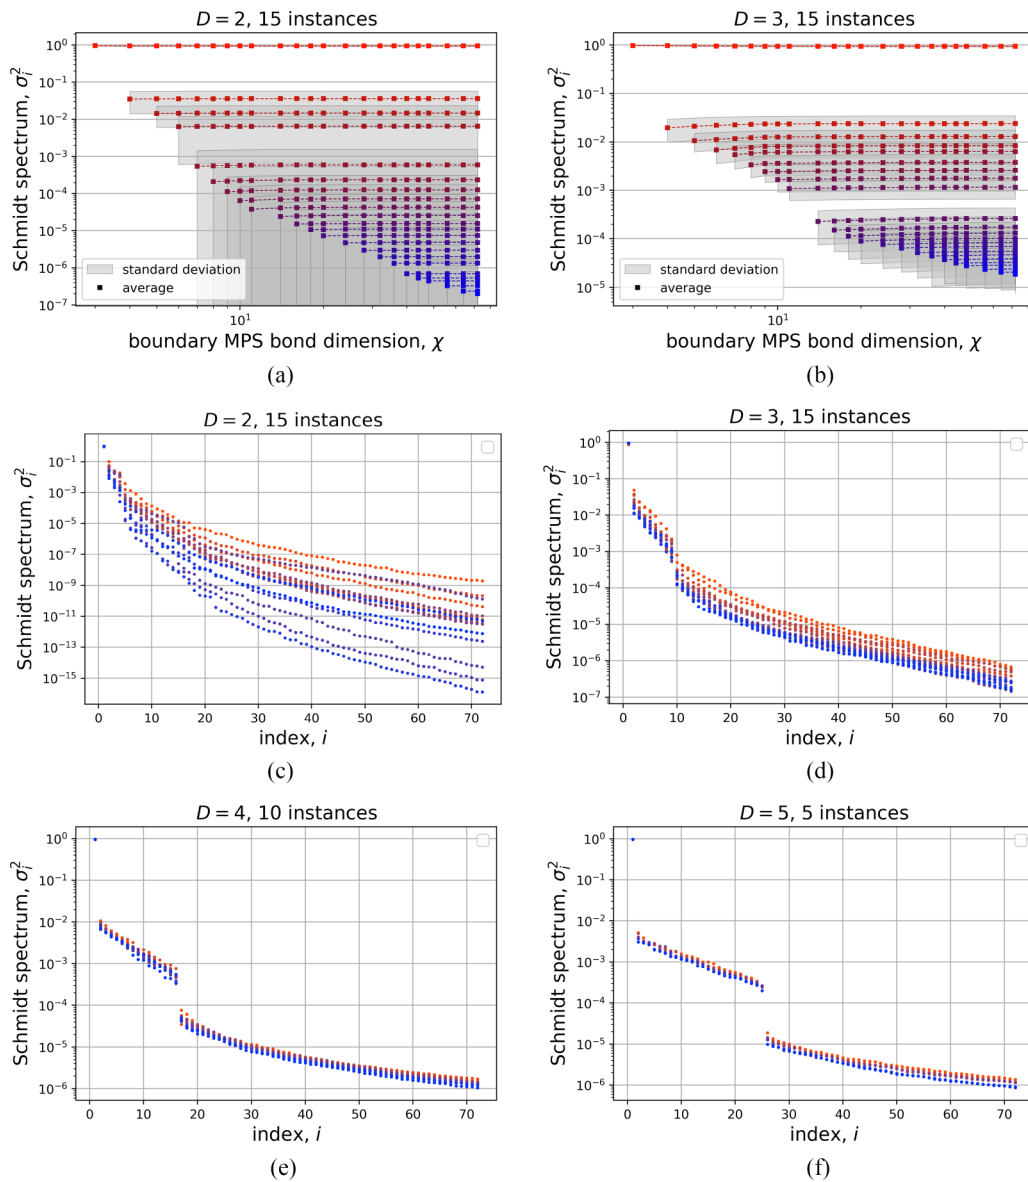
FIG. 9. bMPS average Schmidt entanglement spectrum, several 2d PEPS instances T: [(a) and (b)] $D = 2$ and $D = 3$ (respectively) of the average and standard deviation (over 15 instances) of the Schmidt spectrum of $\tilde{B}(\chi)$ for a range of χ between $\chi_{\min} = 2$, $\chi_{\max} = 72$. The average and standard deviation is calculated for each eigenvalue index. [(c)–(f)] $D = 2, 3, 4, 5$ (respectively) Schmidt spectrum for $\tilde{B}(\chi_{\max} = 72)$ for varying number of instances.

single-site reduced density matrix simply as ϱ. In a PEPS calculation with boundary MPS with bond dimension χ, we obtain an approximation $\varrho(\chi)$ to the exact ϱ. The calculation of $\varrho(\chi)$ is shown in Fig. 6(b), where the dominant left and right eigenvectors $|r_1\rangle$, $\langle l_1|$ are those of the transfer matrix $\mathcal{T}$, depicted in Fig. 8.

The error in $\varrho(\chi)$ due to finite χ in the bMPS depends directly on the square of the Schmidt coefficients of the bMPS, as seen in Fig. 12(a) and therefore on the $n = 1$ (von Neumann) Renyi entropy.

3. iPEPS and boundary MPS: additional numerical results

In this Appendix, we provide additional numerical results for the contraction of 2d infinite PEPS with the boundary MPS formalism.

a. Averaging over several realizations

In the main text we have mostly displayed results for a single instance of a random tensor T for each PEPS bond dimension D. Here, we examine several random instances to build confidence that the reported behavior is typical. Specifically, we explore average properties of $N_s = 15, 15, 10, 5, 1$ instances of clean, random iPEPS with bond dimension $D = 2, 3, 4, 5, 6$, respectively. In all cases, we again observe rapid convergence of the entanglement spectrum (Fig. 9) and Renyi entropies (Fig. 10) with the truncated bMPS bond dimension, χ. Moreover, the variance in these quantities decreases rapidly with increasing PEPS bond dimension D. This suppression qualitatively agrees with the statistical mechanical model, for which fluctuations around the ground state are suppressed at a large D.
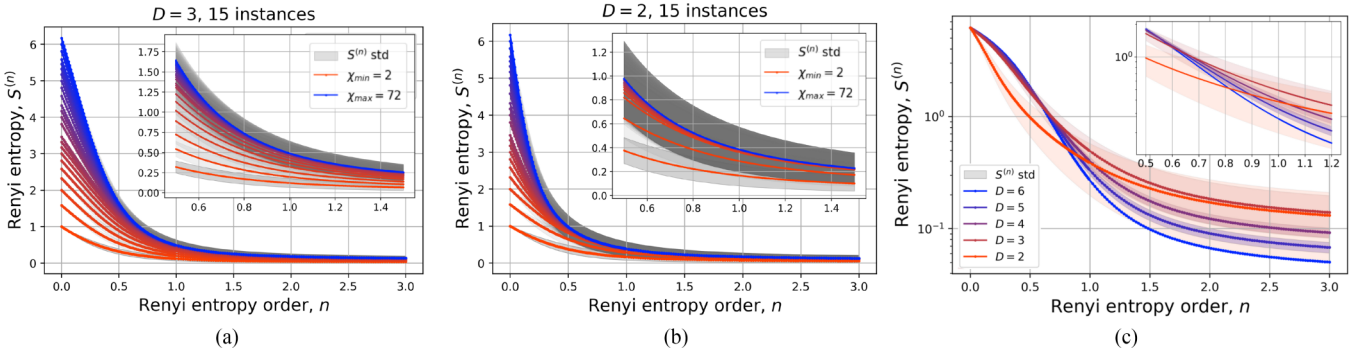
FIG. 10. Renyi entropy order dependence: [(a) and (b)] $D = 2$ and $D = 3$ (respectively) system Renyi entropy as a function of order n, for all simulated bMPS bond dimension $2 \leqslant \chi \leqslant 72$. Average over 15 random initializations of T. (c) Renyi entropy for $\tilde{B}(\chi_{max} = 72)$ as a function of order n for $D = 2, 3, 4, 5, 6$ averaged over 15,15,10,5,1 iterations respectively. We note that the $n = 1$ (von Neumann) entropy is expected to control the error of the truncated bMPS approximation, and that the $n = 0$ entropy simply counts the number of Schmidt weights and trivially saturates to χ_{max}.

b. Fidelities of converged bMPS

In order to quantify how close $|\tilde{\psi}(\chi)\rangle$ is to the exact fixed-point boundary state we calculate the fidelity per site F_s between $|\tilde{\psi}(\chi)\rangle$ and $|\tilde{\psi}(\chi_{max})\rangle$, where we use $|\tilde{\psi}(\chi_{max})\rangle$ as a proxy for $|\tilde{\psi}\rangle$. We see in Fig. 11(a) that this fidelity is very close to 1 already for a very small value of χ, characteristic of weakly entangled states. This small value of χ increases slightly as a function of the 2d PEPS bond dimension D. The sharp transitions in fidelity match the Schmidt spectrum decay in Fig. 1(b). We also calculate the required χ to obtain a certain fidelity per site as a function of D [Fig. 11(b)] and we notice a $\chi \sim D$ relation. This is in contrast with the $\chi \sim D^2$ relation often quoted in the literature for ground states of local Hamiltonians. This discrepancy is likely due to the weakly entangled character of the random PEPS.

c. iPEPS reduce density matrix

In Fig. 1, of the main text we presented the change $\Delta \varrho(\chi)$ in the physical single-site density matrix $\varrho(\chi)$ as a function of χ, where we took $\varrho(\chi_{max} = 72)$ as a reference. We observed that $\Delta \varrho(\chi)$ is exponentially suppressed with χ. We expect this result to apply also to the density matrix $\varrho(\chi)$ on a larger local region. As an example, Fig. 12(b) shows $\Delta \varrho(\chi)$ for $\varrho(\chi)$ for a

region of two contiguous sites, where indeed we again observe exponential suppression with χ.

d. Correlation length

The correlation length that we obtain from the converged bMPS at maximum bond dimension are low: averaging at order $\sim$1 for all bond dimensions D [Fig. 13(b)]. The convergence of the correlation length as a function of bond dimension χ [Fig. 13(a)] also matches the gaps and exponential decay in the spectrum as per Fig. 1. This also implies that for relatively low bond dimensions we obtain correlation lengths that are comparable to our best estimate given by χ_{max}.

e. bMPS entanglement entropy for the 2d tensor network corresponding to the overlap of two iPEPS

We have numerically and analytically confirmed that for 2d random PEPS the approximate computation of the norm $\mathcal{N}$ and expectation values of local observables can be performed efficiently. However, other types of observables, such as individual wave-function components, $\langle s_1 \ldots s_N | \Psi \rangle$, or overlaps between distinct random PEPS, $\langle \Psi' | \Psi \rangle$, are expected to be hard since they involve a 1d transfer matrix map $T T'^*$, where $T \neq T'$, which is not positive.

We consider the computation of the overlap of two distinct 2d random PEPS $\langle \Psi | \Psi' \rangle$ given by tensors T and T', where we
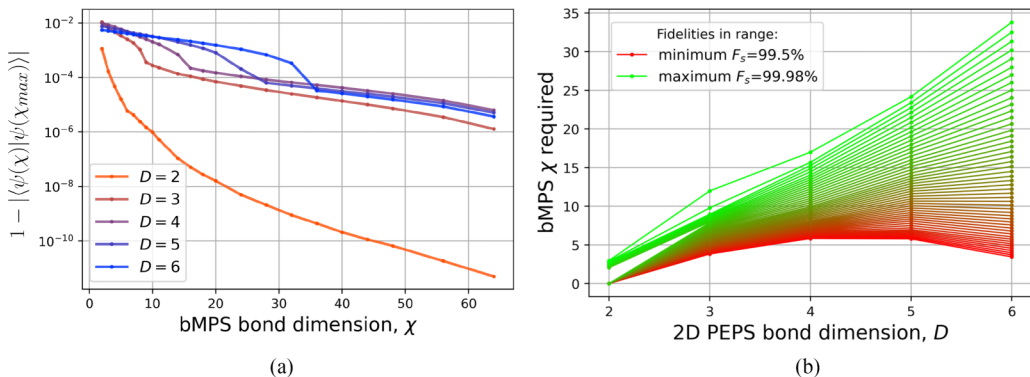

FIG. 11. bMPS fidelity numerical results. (a) Convergence of the fidelity per site of the largest simulated bMPS network ($\chi_{max} = 72$) with all other smaller systems. (b) Bond dimension required for the bMPS to obtain a certain fidelity per site with the largest simulated boundary.
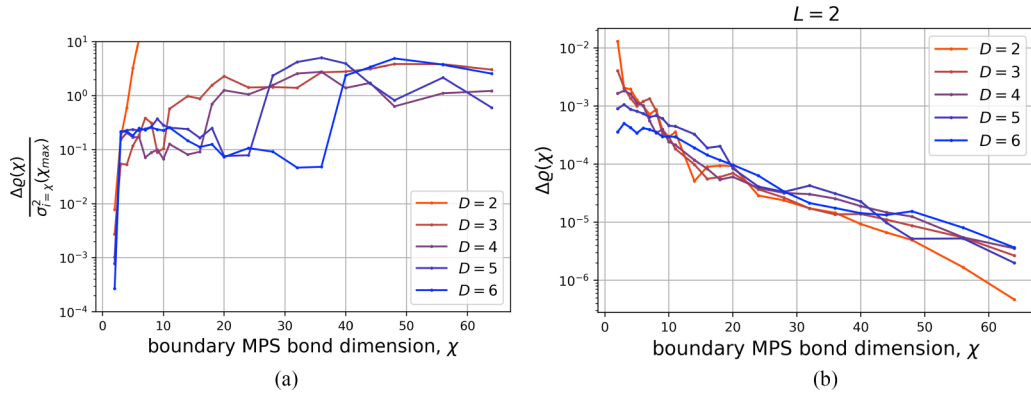
FIG. 12. 2d PEPS reduced density matrix. (a) Largest eigenvalue of the difference between the largest and smallest bMPS $\rho(L)$ divided by the squared Schmidt spectrum of the boundary for χ_{max}. This plot shows a constant dependence, indicating that the order-1 Renyi entropy is to be considered when analyzing observable accuracy of the 2d PEPS contraction. (b) Change $\Delta\varrho(\chi)$ in the physical $L = 2$ density matrix $\varrho(\chi)$ as a function of χ

build T' from T with noise tuned by a parameter $\eta \in [0, 1]$, such that $T' = (1 - \eta)T + \eta\mathcal{C}_\mathcal{N}$, where $\mathcal{C}_\mathcal{N}$ is the normal complex ditribution with zero mean and standard deviation 1. In the results in Fig. 14, we observe an increase in the Schmidt entanglement entropy at the boundary for the contraction evolution of the overlap $\langle\Psi|\Psi'\rangle$, where for large η, the entanglement saturates to a constant value given by the bond dimension χ of the bMPS.

APPENDIX B: STABILIZER PEPS FORMALISM AND ADDITIONAL NUMERICAL RESULTS

In this Appendix, we provide details about the stabilizer PEPS formalism and its simulation.

A stabilizer PEPS is a PEPS whose composing tensors $\{T_{[\mathbf{r}]}\}$ are stabilizer tensors, i.e., each of the states defined through

$$|T_{[\mathbf{r}]}\rangle = \sum_{i,j,k,l,s} (T_{[\mathbf{r}]})_{s,ijkl}|s, ijkl\rangle, \tag{B1}$$

is a quantum stabilizer state over p-qudits, for some prime number p. In the expression above, i, j, k, l are virtual bond indices, and s is the physical bond index. Each bond represents a Hilbert space containing an integer number of p-qudits, thus

both physical and virtual bond dimensions need to be some integer power of p: $D = p^{k_D}, d = p^{k_d}$.

The contraction of two stabilizer tensors can be simulated efficiently by performing several Bell measurements on the contracted bonds, as detailed below. Let us assume $T^1_{ab_1}$ and $T^2_{b_2c}$ are two stabilizer tensors where the two b indices are of the same dimension p^k. Let T^{12}_{ac} be the tensor obtained by contacting the two tensors over the b indices:

$$T^{12}_{ac} := \sum_b T^1_{ab}T^2_{bc}. \tag{B2}$$

The contraction can be realized by performing several (forced) Bell measurements on qudits within b_1 and b_2 [45]. If we label the qudits in the b_1 with $\{1_1, \ldots, k_1\}$, and those in the b_2 with $\{1_2, \ldots, k_2\}$, then

$$|T^{12}\rangle \propto \prod_{i=1}^{k} \left(\mathcal{M}[X_{i_1}X_{i_2}]\mathcal{M}[Z_{i_1}Z_{i_2}^{-1}]\right)|T^1 \otimes T^2\rangle, \tag{B3}$$

where $\mathcal{M}[P]$ is the projector to the $+1$ subspace of the Pauli operator P. Using the Gottesman-Knill theorem [46], one can show that the state $|T^{12}\rangle$ is still a stabilizer state. Further, its stabilizers can be obtained from those of T^1 and T^2 with
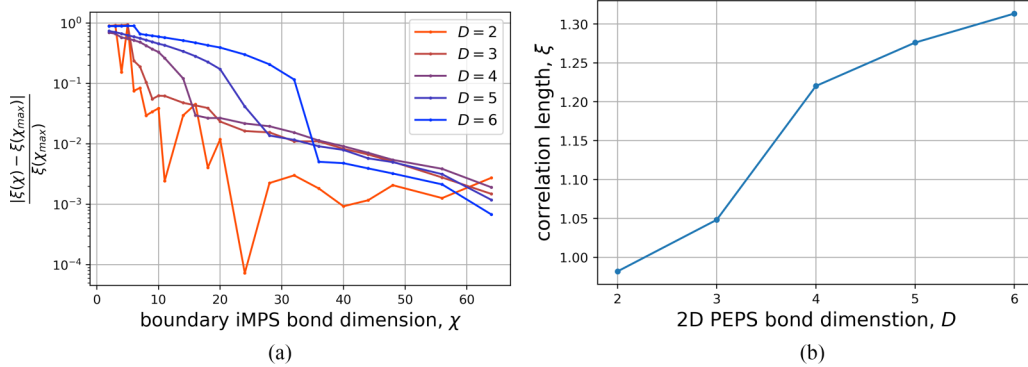

FIG. 13. bMPS correlation length numerical results. (a) Convergence of the correlation length as a function of bMPS bond dimension χ and for varying 2d PEPS sizes D. (b) Average correlation length for $\tilde{B}(\chi_{max} = 72)$ as a function of D for $D = 2, 3, 4, 5, 6$, averaged over 15, 15, 10, 5, 1 initialization, respectively.
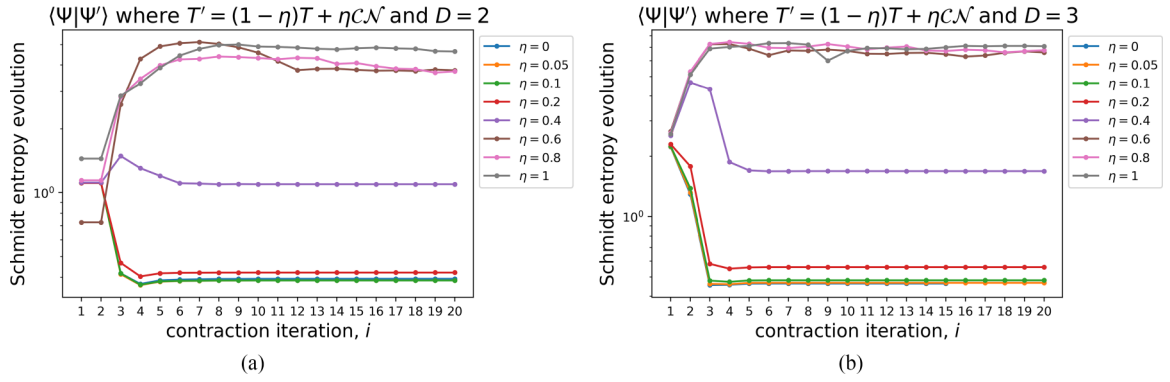
FIG. 14. Global vs local observables: [(a) and (b)] show the evolution of the entanglement entropy at the boundary MPS for the overlap $\langle \Psi | \Psi' \rangle$, where both 2d PEPS have bond dimension $D = 2$ and $D = 3$ (respectively) and the bMPS has bond dimension $\chi = 35$.

$O(\ln^3 D_{\text{total}})$ time complexity, where $D_{\text{total}} = \dim a \cdot \dim c \cdot (\dim b)^2$ is the product of all the open bonds' dimensions.

It is worth noting that the complexity of contracting stabilizer tensors is independent of the entanglement property of the underlining states. Thus for a given stabilizer PEPS, we are able to compute the boundary state $|\psi(t)\rangle\rangle$'s evolution *exactly* without any truncation and study its entanglement properties. The latter can be computed from the state's stabilizers using the algorithms introduced in Ref. [47].

We consider both disordered and clean (i.e., translational invariant) stabilizer PEPS. In both cases the PEPS is finite and takes periodic boundary conditions along the x direction, thus the boundary state $|\psi(y)\rangle\rangle$ is also finite and has the periodic boundary condition. To reach the large-D regime, we take $p = 173$. The simulated von-Neumann entropy $S_A(y)^3$ of $|\psi(y)\rangle\rangle$ is presented in Fig. 3 in the main text as well as Fig. 15.

We start by focusing on the disordered case, where each unit stabilizer tensor $T_{\mathbf{r}}$ is sampled independently. The simulation shows that at any given layer number y and when $|A|$ is far from 1 or L, the von Neumann entropy $S_A(y)$ takes a

---

[3]For a stabilizer state (tensor) and any bipartition of bonds, the nonzero singular values associated with the bipartition are all of the same value. Thus the Rényi entanglement entropy across the bipartition is independent of the Rényi index α, in particular, it equals the von Neumann entropy ($\alpha = 1$).

constant value that is independent of $|A|$ or the system size L (Fig. 15, left). Further, the constant value first increases linearly and then drops with the increase of the number of contracted layers y. Both the peak value and the turning point are dependent on the bond dimensions D and d (Fig. 15, mid). The two plots together suggest that the boundary state $|\psi(t)\rangle\rangle$ is at most area-law entangled at any time (layer number) y. Further, the simulated entanglement barrier's dependence on (d, D) matches with the prediction of the stat-mech mapping in Eq. (5), as is shown in Fig. 3 in the main text.

Next we come to the clean PEPS case, where all unit tensors $T_{[r]}$ are identical and taken to be a randomly sampled stabilizer tensor. The simulation results suggest that the behavior of $S_A(y)$ is almost identical to that in the disordered case. Namely, the $|\psi(y)\rangle\rangle$ is also area-law entangled in the clean case, with an area-law value following the expression Eq. (5) [Figs. 15 (right) and 3 (right)].

APPENDIX C: DETAILS OF STATISTICAL-MECHANICS MAPPING

In this Appendix, we review the derivation of the stat-mech description of entanglement features of random PEPS.

In the Gaussian random PEPS ensemble, the tensor T_r for each site r is chosen independently and identically distributed from a Gaussian distribution:

$$\mathbb{E}\left[(T_{[r']}^{*})_{i'j'k'l'}^{s'}(T_{[r]})_{ijkl}^{s}\right] = \delta_{ss'}\delta_{ii'}\delta_{jj'}\delta_{kk'}\delta_{ll'}\delta_{r,r'}, \quad \text{(C1)}$$
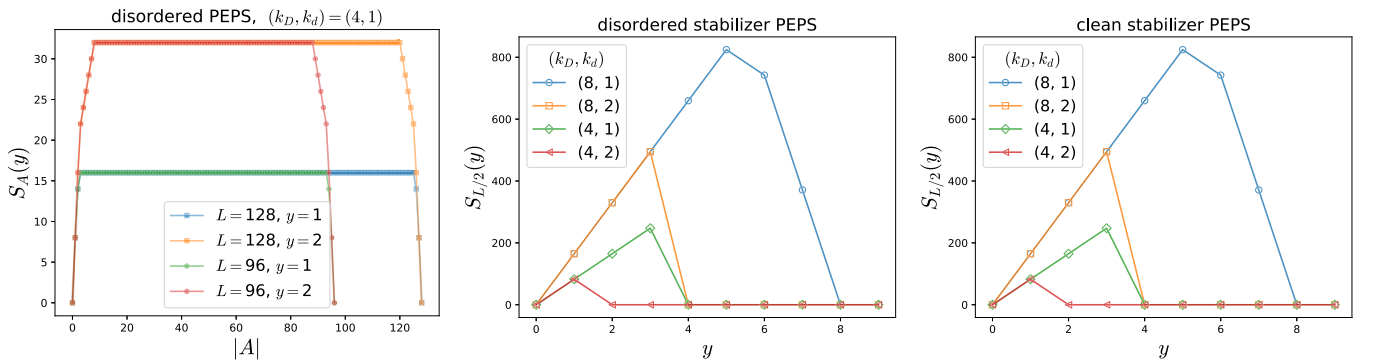

FIG. 15. Further numerical results for stabilizer PEPS. (Left) Profile of $|\psi(y)\rangle\rangle$'s entanglement entropy in a disordered PEPS. The L- and $|A|$- independent plateau suggests area-law entanglement of $|\psi(y)\rangle\rangle$. (Middle) The simulated "entanglement barrier" for various choices of D and d in the disordered PEPS. (Right) The simulated "entanglement barrier" in the clean PEPS.

where $\mathbb{E}[\dots]$ denotes averaging over the ensemble, $s = 1, \dots, d$ is the physical index, and $i, j, k, l = 1, \dots, D$ are bond indices, and r label sites of the lattice. In the following, we drop the indices s, i, j, k, l on the tensors.

1. Mapping RTNs to replica magnets

Consider the tensor network contraction to compute the norm of the PEPS: $\mathcal{N} = \langle \Psi | \Psi \rangle$ using the MPS method outlined in Sec. A above. Denote the (un-normalized) density matrix of the evolved boundary state as $\rho(y) = |\psi(y)\rangle\rangle\langle\langle\psi(y)|$. Our aim will be to compute the evolution of the ensemble-averaged Renyi entanglement entropy of a region A of the evolved boundary state, $|\psi(y)\rangle\rangle$.

$$S_A^{(n)} = \frac{1}{1-n} \ln \frac{\mathrm{tr}\rho_A^n}{(\mathrm{tr}\rho)^n}, \tag{C2}$$

where and ρ_A is the reduced density matrix in some contiguous interval A of size L_A obtained from tracing out the complement of A in $\rho = |\psi\rangle\rangle\langle\langle\psi|$.

Since the wave function $|\psi\rangle\rangle$ is not necessarily normalized, so the denominator in Equation (C2) is crucial to obtain a meaningful entanglement entropy. Directly computing the disorder average of this nonlinear quantity is challenging. To avoid this difficulty, we employ a standard replica trick based on the identity

$$\ln \mathrm{tr}\rho_A^n = \lim_{m \to 0} \left((\mathrm{tr}\rho_A^n)^m - 1 \right) / m. \tag{C3}$$

This allows us to express the disorder average of Eq. (C2) as

$$\mathbb{E}[S_A^{(n)}] = \frac{1}{n-1} \lim_{m \to 0} \frac{1}{m} (\mathcal{F}_A - \mathcal{F}_0), \tag{C4}$$

with $\mathcal{F}_{A,0} = -\ln \mathcal{Z}_{A,0}$ and $\mathcal{Z}_0 \equiv \mathbb{E}[(\mathrm{tr}\rho^n)^m]$, $\mathcal{Z}_A \equiv \mathbb{E}[(\mathrm{tr}\rho_A^n)^m]$. Using this identity, the calculation of the Renyi entropies reduces to computing $\mathcal{Z}_0$ and $\mathcal{Z}_A$, and to evaluate the replica limit (C4).

When m and n are integers, the averages in $\mathcal{Z}_0$ and $\mathcal{Z}_A$ can be evaluated analytically using Wick's theorem. One can then express the partition functions $\mathcal{Z}_A$ and $\mathcal{Z}_0$ in terms of a classical statistical mechanics model, whose degrees of freedom are *permutations* labeling different Wick contractions at each vertex of the tensor networks: at each site, each tensor T_r must be paired with a T_r^* possibly belonging to a different replica. Let $Q = nm$ be the number of copies of ρ. Then, the partition function $\mathcal{Z}$ involves computing quantities like $\mathbb{E}[\rho^{\otimes Q}]$. Note that $|\psi\rangle\rangle$ contains both T_r with T_r^* at each site r. Then, in the replicated theory there are $2Q$ copies of T, and $2Q$ copies of T^* for each site, which we label, T_r^α with a replica index $\alpha = 1, \dots, 2Q$. We adopt the following ordering for the $2Q$ copies:

$$\{1, \bar{1}, 2, \bar{2}, \dots, Q, \bar{Q}\}, \tag{C5}$$

where k labels the state ("ket") in replica Q, and $\bar{k}$ denotes the dual state ("bra") in replica Q. To label permutations we use cycle notation, for example $(124)(35) \in S_6$ denotes the permutation $123456 \to 245136$, i.e., with separate cyclic permutations of elements (124) and of elements (35) (for convenience, we only list the cycles with more than one element). We will also need to define the cycle counting function

$$C(g, h) \equiv C(g^{-1}h) = \# \text{ of cycles in } g^{-1}h, \tag{C6}$$

where $C(g)$ also includes single-element cycles that are not listed explicitly in our notation, e.g., for the above example, $C[(124)(35)] = 3$.

According to Wick's theorem, upon averaging over the Gaussian random tensors, a nonzero contribution is obtained only if each tensor T_r^α is paired with a permuted copy $(T_r^{g_{[r]}(\alpha)})^*$, where $g_{[r]} \in S_{2Q}$ labels a permutation of the replicas, and S_{2Q} is the symmetric group on $2Q$ elements.

The partition function $\mathcal{Z}$ corresponding to the tensor network contraction, can then be written as a sum over replica-permutation "spins" for each site:

$$\mathcal{Z} = \sum_{\{g_{[r]}\}} W[\{g_{[r]}\}], \tag{C7}$$

where W is the weight of the Wick contraction for the corresponding spin configuration. The weight can be computed analytically for each bond in the tensor network. There are three distinct types of contractions to consider.

(1) Bulk bonds connecting different nearest-neighbor nodes i and j with permutation "spins" $g_{[r]}$ and $g_{[r']}$, and bond dimension D. The same contraction occurs in the $2Q$ layers of the replicated tensor network. Since Wick contractions force indices to be the same, the resulting weight is equal to D:

$$D^{C(g_{[r]}, g_{[r']})} = e^{(\ln D)C(g_{[r]}, g_{[r']})}, \tag{C8}$$

on each link of the square lattice, since the number of independent bond indices is equal to $C(g)$, the number of cycles in the permutation g. This is most easily seen by a graphical representation: each cycle in $g_{[r]}^{-1}g_{[r']}$ leads to a "loop" where indices have to be the same, with weight $\sum_{\alpha=1}^D \delta_{\alpha\alpha} = D$. Interpreting this positive weight as a Boltzmann weight, this terms leads to a ferromagnetic interaction (favoring $g_{[r]} = g_{r'}$ for neighboring i, j to maximize the number of cycles to $2Q$) with interaction strength $\ln D$. This Boltzmann weight has a left/right symmetry $(S_{2Q} \times S_{2Q}) \rtimes \mathbb{Z}_2$ (where the extra $\mathbb{Z}_2$ symmetry corresponds to $g \to g^{-1}$).

(2) Bulk contraction of T_r with T_r^* along the physical leg with dimension d. This contraction can be implemented by adding a site with fixed permutation equal to identity $e = ()$: in each replica k, we pair k with itself (corresponding to gluing T with T^* in the ket), and $\bar{k}$ with itself (corresponding to gluing T with T^* in the bra). This leads to a factor

$$d^{C(eg_{[r]})} = e^{(\ln d)C(g_{[r]})}, \tag{C9}$$

on each site. This can be seen as a $S_{2Q} \times S_{2Q} \to S_{2Q}$ symmetry-breaking field favoring the identity permutation. This bulk field on every site prevents any phase transition by creating an energy cost for domains of spins with $g_{[r]} \neq e$ that scales as the volume of the domain.

(3) Boundary contractions at the top layer to compute $\mathcal{Z}_0 \equiv \overline{(\mathrm{tr}\rho^n)^m}$, $\mathcal{Z}_A \equiv \overline{(\mathrm{tr}\rho_A^n)^m}$. At the top layers we have dangling legs, that should be contracted to implement the trace and partial trace operations to compute entanglement. In $\mathcal{Z}_0$, we want to compute $\mathrm{tr}\rho$ in each replica. This means that in each replica (and at each site at the boundary), we want to contract T_r (respectively T_r^*) in the ket with T_r^* (respectively T_r) in the bra. In our language, this corresponds to the permutation (Fig. 16)
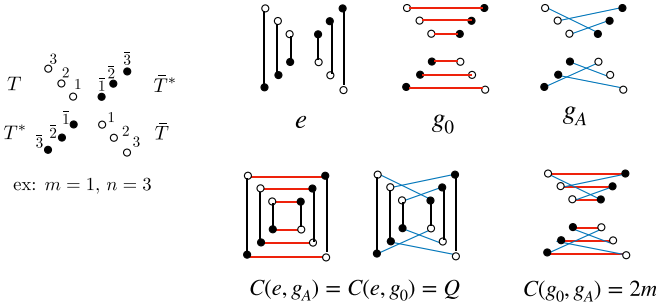
$$g_0 = (1\bar{1})(2\bar{2}) \dots (Q\bar{Q}), \tag{C10}$$

FIG. 16. Graphical notation for permutations and cycle counting. (Left) In each of the $Q = mn$ replicas of $\rho = |\psi\rangle\langle\psi|$, there are $2Q$ T tensors, and $2Q$ T^* tensors. We label T, T^* tensors from the bra $\langle\psi|$ with an overbar. It is convenient to draw the T and T^*'s in m groups of n (shown here for $m = 1$). Averaging over tensors forces a wick contraction between T_i and $T^*_{g(i)}$ where $g \in S_{2Q}$. (Top right) The three permutations corresponding to the bulk ($e = $ identity) and boundary (g_A, g_0) fields. (Bottom right) Graphical calculation of cycle counting for various fields, $C(\sigma)$ is given by counting the number of independent loops in the picture (note, for $m > 1$, not shown, there would simply be m independent copies of this picture).

Note that this permutation is *not* identity, it has Q cycles while $e = ()$ has $2Q$ cycles. At the end of each leg, we fixed the permutation to g_0:

$$D^{C(g_0^{-1}g_{[r]})} = e^{(\ln D)C(g_0 g_{[r]})}, \tag{C11}$$

for $i = 1, \ldots, L$ at the top boundary. To implement the partial trace in $\mathcal{Z}_A \equiv \overline{(\text{tr}\rho_A^n)^m}$, we fixed the permutation to g_0 if i is in $\overline{A}$, and to g_A is $r \in A$, with

$$g_A = ((1\overline{2})(2\overline{3})\ldots(n\overline{1}))^{\otimes m}. \tag{C12}$$

Assembling these ingredients results in the effective Hamiltonian of the main text.

2. Comparison to related stat-mech models

A nearly identical replica-spin model was derived in Ref. [18] for holographic random tensor network states (rTNS), i.e., whose tensors had physical legs only at the boundary, and only virtual bond legs in the bulk. In that work, the key difference was that the holographic rTNS did not have positive tensors. As a result the permutation spins were S_Q- rather than S_{2Q}-valued, and the bulk retained the S_Q symmetry since there was no field along e. This crucial difference led to two possible phases of the stat-mech model: a disordered (paramagnetic) phase at weak coupling (low-D) in which the permutation spins are short-range correlated, and an ordered (ferromagnetic) phase at strong coupling (large-D) in which the S_Q symmetry was spontaneously broken and the permutation spins have long range order. In the ordered phase, domain walls had a nonvanishing surface-tension, resulting in an extensive free energy for the boundary twist in the entanglement region size, resulting in volume-law entanglement scaling. By contrast, in the disordered phase, there is only a local free-energy cost at the edge of the boundary domains, resulting in area-law entanglement.

Coming back to the stat-mech model for the PEPS norm computation: the key difference is that the tensors are completely positive, i.e., are composite tensors made up of T and T^* with physical legs contracted. This results in a bulk field along the identity (e) permutation that explicitly breaks the S_{2Q} symmetry. A similar statistical mechanics model emerges in the context of random quantum channels [25]. In that language, bond dimension corresponds to physical dimension in the channel, and our physical dimension d maps to the strength of channel (effectively the number of Kraus operators).

Intuitively, this explicit symmetry breaking destroys the ordering transition, such that the entire phase diagram is effectively ordered (in the sense that domain walls have a nonzero surface tension). Naively, one might expect this to result in a volume-law entanglement throughout for any D. However, as we show next, there is an exact cancellation of the volume-law contribution to the free energy with twisted boundary conditions, generically resulting in area-law entanglement for the evolved boundary state.

3. Minimum cut picture of random PEPS contraction

At a large D, the permutation spins are strongly locked to each other by their ferromagnetic interactions, and pinned to the bulk e fields. Here, domain walls have a nonzero line tension, and we can approximately compute the domain wall free energy for the entanglement entropy, by minimizing this line tension.

Let us focus on the thermodynamic limit, $L, y \to \infty$, and on half-system entanglement. The statistical mechanics model has a bulk symmetry-breaking field that prevents an entanglement phase transition which would be associated with a spontaneous breaking of S_{2Q} symmetry. Specifically, the e field produces an energy cost for domains with $g_{[r]} = e$ that scales like the volume of the domain. This field favors the identity permutation e in the bulk (with fluctuations suppressed if $D \gg 1$), while the boundary fields favor g_0 or g_A. However, the energy cost of the domain walls between those permutations and e are the same, since g_A and g_0 each have $Q = nm$ cycles. Therefore $\mathcal{F}_A$ and $\mathcal{F}_0$ each contain an extensive term $L \ln D$, but importantly this extensive term is the same for both $\mathcal{F}_{A,0}$ and cancels in the difference. This cancellation can be traced back to the $S_Q \times S_Q$ symmetry of the e field in the bulk, and that g_0 and g_A differ by a transformation in this symmetry group so that the two types of boundary conditions are locally equivalent. Consequently, the only difference between A and 0 arises from local energy cost associated with the domain wall between g_0 and g_A BCs, which in 2d has constant size independent of L. In general, we thus have

$$S_A^{(n)} \sim \text{constant}. \tag{C13}$$

That is, the top boundary is always area law. In fact, as $D \to \infty$, both partition functions are dominated by the ground-state configuration where all spins are $g_{[r]} = e$, and we have $\mathcal{F}_A - \mathcal{F}_0 = 0$ corresponding to a disentangled state. We note, in passing, that an identical calculation for a 3D PEPS shows that the operator entanglement would scale linearly in the system size. It is plausible that a boundary 2d PEPS, which can account for such linear scaling of operator entanglement, would again enable an efficient contraction.

4. Fluctuation corrections

As we now briefly discuss, fluctuation corrections to the $D = \infty$ limit can be viewed as an expansion in dilute gas of spin flips on top of the e-polarized ground state.

$$
E_1(g) = \begin{cases} (4J + h)c(g, e) & \text{bulk} \\ (3J + h)c(g, e) + h[c(g, g_{A/0}) - c(e, g_{A,0})] & \text{r} \in A/\bar{A} \text{ boundary} \end{cases}. \tag{C14}
$$

The lowest-energy spin flips correspond to transposing a single pair of replicas: $g = (ab)$, which have bulk energy: $E_1((ab))_{\text{bulk}} = 4J + h$. There are $D_1^{\text{bulk}} = Q(2Q - 1)$ different single-transposes, leading to a corresponding degeneracy of the single SF excitations in the bulk. Near an $\bar{A}$ boundary, the cheapest spin flip is is of the form $(k\bar{k})$ for some $k = 1, \ldots, 2Q$, and costs energy $E_1((k\bar{k}))_{\bar{A}-\text{bdry}} = 2J + h$, and degeneracy $D_1^{\text{edge}} = Q$. The minimal-energy spin flips and corresponding degeneracy near A boundary are related to those near the $\bar{A}$ boundary by the symmetry generator: $g_A g_0^{-1}$, which commutes with the bulk e fields. At a large D, we can use these excitations to approximate the free energy by a dilute gas of spin-flip excitations. This expansion is however subtle in the replica limit, as permutations with a number of cycles proportional to $m \to 0$ become dominant in the replica limit. While this caveat prevents us from systematically computing the free energy in a controllable way, this simple counting of low energy excitations predicts that the area-law coefficient scales as

$$
\mathcal{F}_A - \mathcal{F}_0 = -\ln \frac{\mathcal{Z}_A}{\mathcal{Z}_0} \approx D^{-2}d^{-2} + \mathcal{O}(D^{-3}). \tag{C15}
$$

Though we are unable to explore a large enough range of D with sufficient precision in the iPEPS numerics to test this asymptotic prediction in detail, we note that the large-D suppression in sample-to-sample variance of entanglement features observed in the iPEPS numerics is in qualitative agreement with the suppression of fluctuation contributions to the stat-mech model at a large D.

5. Correlation length of random PEPS

The stat-mech mapping also enables one to estimate the correlation length-scale for observables in random PEPS. Namely, consider computing the typical amplitude of a correlation function:

$$
\ln C_{\text{typ}} \equiv \mathbb{E}\left[\ln\left[\frac{\langle\Psi|O_1 O_2|\Psi\rangle\langle\Psi|\Psi\rangle}{\langle\Psi|O_1|\Psi\rangle\langle\Psi|O_2|\Psi\rangle}\right]\right], \tag{C16}
$$

where $O_{1,2}$ are local observables on sites 1,2, and we have explicitly normalized the wave-function, and also divided by the (normalized) one-point correlators: $(\frac{\langle\Psi|O_{1,2}\|\Psi\rangle}{\langle\Psi|\Psi\rangle})$ to remove dependence on the operator norm of $O_{1,2}$.

For concreteness, and without loss of generality, let us specialize to the case where $O_{[r]} = |s\rangle_{[r]}\langle s|$ is a projector onto physical state $|s\rangle$ at site r (and identity elsewhere). Generic observables can be written as linear combinations of such projectors (up to a basis transformation that can be absorbed into the randomly drawn tensor on site $[r]$).

For a finite number of replicas, the minimal-energy excitations are single spin flips (1SF's), where we replace $g_{[r]} : e \to \sigma \neq e$ at some site r. Denoting $J = \ln D$ and $h = \ln d$, and $c(g, g') = C(g, g') - C(e, e)$, the 1SF costs energy:

One can evaluate the average of the logarithm in C_{typ} via a replica trick as outlined above for the bMPS entanglement. The result is that

$$
C_{\text{typ}} = \exp\left[-\left(F_{O_1, O_2} - F_{O_1} - F_{O_2} + F_0\right)\right] = \frac{\mathcal{Z}_{O_1 O_2}\mathcal{Z}_0}{\mathcal{Z}_{O_1}\mathcal{Z}_{O_2}}, \tag{C17}
$$

where F_0 is the free energy associated with the stat-mech Hamiltonian (3), $F_{O_1, O_2, \ldots, O_k}$ is the same free energy except with the projectors inserted at sites $1, 2, \ldots, k$, and $\mathcal{Z} = e^{-F}$ is the associated partition function. The projectors restrict the sum over the physical index values to s, which is equivalent to removing the replica symmetry breaking: e field on that site. Equivalently, $F_{O_1, O_2, \ldots, O_k}$ is given by the free energy of the Hamiltonian discussed in the main text, but perturbed by a term: $\Delta H = +\ln d \sum_{[r]=1, \ldots, k} C(e, g_r)$.

At a large D, we can estimate the leading contribution to C_{typ} as follows. The leading contribution to the stat-mech partition function is from a uniformly e-polarized replica-spin configuration. Fluctuations about this come in the form of small domains of non-e polarized spins. By inspection, only domains that include both sites 1,2 make a noncancelling contribution to the ratio in (C17). At a large D, this contribution is dominated by the smallest such spanning domain, which is a line of flipped replica spins, $g_i \neq e$, along a short path connecting points 1 and 2. This domain has a line tension $\mathcal{F} \approx (\ln D^4 d)r_{12}$, where r_{12} is the length of the shortest path through the network connecting points 1,2. This contributes exponentially decaying correlations:

$$
C_{\text{typ}} \approx e^{-r_{12}/\xi}, \tag{C18}
$$

with characteristic correlation length: $\xi \approx (\ln dD^4)^{-1}$. Notice that the correlation length decreases with increasing D. However, note that random PEPS states at a large D *are not* close to product states, but, in fact saturate the maximal entanglement allowed for the given bond-dimension PEPS.

This shows that large-D random square lattice PEPS actually have rather short-range correlations, in accordance with previous studies [21], and our numerical observations for clean random iPEPS.

APPENDIX D: RANDOM PEPS VERSUS PHYSICALLY RELEVANT GROUND STATES

An important question is to what extent the results presenting in this paper, strongly indicating that random PEPS can be efficiently approximately contracted, can be extended to PEPS representing ground states of physically relevant Hamiltonians, e.g., in the context of condensed matter, materials
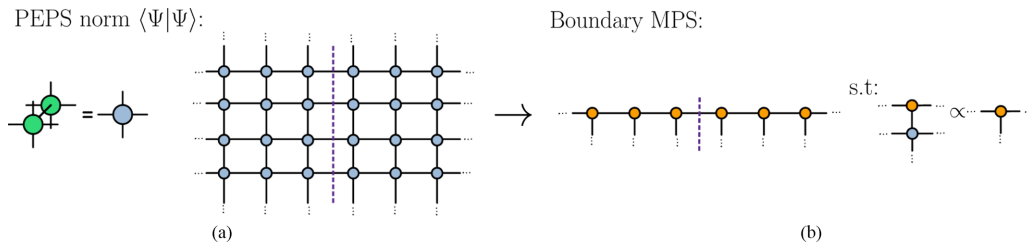
FIG. 17. (a) Square lattice of the 2D PEPS norm. (b) Boundary MPS obtained from the contraction of the norm in (a).

science and quantum chemistry. That is, can our results shed some light into the performance of PEPS algorithms when simulating such systems? Here we restrict our considerations to two-dimensional ground states that obey the entanglement area law (2d ground states that violate the entanglement area law, such as ground states in the presence of a 1d Fermi surface, are expected to be harder to contract).

We have seen below that random PEPS have a very short correlation length ξ on the order of one lattice site or less. In contrast, the correlation length ξ in a physically relevant ground state can be arbitrarily large (for instance, the correlation length diverges as we approach a quantum critical point). Relatedly, we have numerically seen that the boundary MPS for a random PEPS has very limited amount of entanglement whereas, as PEPS practitioners have learned over the last 15 years, the entanglement in the boundary MPS for a physically relevant 2d ground state can again be arbitrarily large (even in those cases where the boundary MPS obeys an area law).

We have therefore identified two structural differences between random PEPS and physically relevant ground states, namely, differences in correlation lengths ξ and in the amount of boundary MPS entanglement. How fundamental are these structural differences? Based on experience with renormalization group, random-circuit dynamics, and random matrix theory it is tempting to conjecture that the large-D random PEPS might represent a sort of coarse-grained "fixed-point" representation of physically relevant ground states. However, below we will see that while coarse-graining the PEPS for a physically relevant ground state would indeed effectively remove the difference in correlation length, it would not change the difference in boundary MPS entanglement. Since boundary MPS entanglement determines the computation cost in approximate PEPS contractions, we cannot conclude that our results for random PEPS apply to such PEPS.

That is not to say that the stat-mech approach used in this paper to successfully characterize the boundary MPS entanglement for random PEPS is restricted to studying states with a short correlation length ξ. On the contrary, as discussed below, we will see that the same approach can be used for 2d random tensor network states (which are not 2d PEPS) with arbitrarily large correlation length ξ.

1. Coarse-graining by blocking tensors

First, note that, any finitely correlated PEPS, i.e., with finite correlation length, ξ, can be transformed into a PEPS with shorter correlation length $\xi' \sim 1$ by "blocking" together sites in $\xi \times \xi$ blocks of physical sites. This blocking adds constant

overhead to the bond-dimension of each tensor, $D' \sim D^\xi$. While this cost may be severe in practice, from an asymptotic complexity standpoint, it is merely a constant overhead. This argument suggests that one can perhaps think of a random PEPS as reflecting a block-spin renormalization group (RG) style coarse-graining of a physical PEPS with longer-range correlations.

However, the following observation suggests that there is no connection between a coarse-grained PEPS for physical models with large ξ, and a random PEPS with bond-dimension D'. Years of numerical experience [5,8–13] show that ground states with large correlation length have corresponding large entanglement both for the physical state, and the bMPS for its norm and correlation functions. While blocking reduces the correlation length, it does not change the entanglement spectrum of the bMPS (see Figs. 17 and 18). Hence, for physical states, the bMPS entanglement should grow with ξ, whereas for random PEPS with bond dimension $D' = D^\xi$, the method of Appendix C 4 above predicts bMPS entanglement decreasing exponentially with ξ as $\sim 1/D^{2\xi}$. This argument shows that a large D random PEPS do not have the correct entanglement structure to capture block coarse-grainings of long-range correlated states encountered in simulation of physical systems.

However, in standard renormalization group approaches, coarse-graining does not simply involve merely blocking sites together, but of hierarchically decomposing the state via layers of coarse-graining steps that act on different distance scales. Inspired by this, in the next section, we construct a class of hierarchical random tensor network states that have arbitrarily long correlation length ξ, and which have bMPS entanglement that grows with ξ in a manner qualitatively consistent with that found in simulations of physical systems (though it remains an open question whether variants of such hierarchical tensors
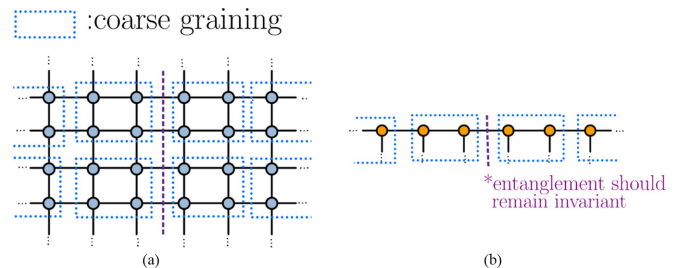

FIG. 18. (a) Coarse grained square lattice of the 2D PEPS norm. (b) Boundary MPS obtained from the contraction of the coarse grained norm.
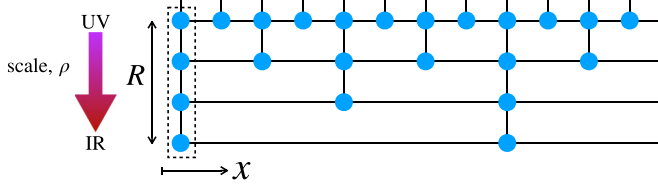
FIG. 19. Random tensor networks with arbitrarily long correlation length. Schematic of 1*d* gMERA (with obvious generalization to higher-dimensional versions) which has arbitrarily long correlation length $\xi \sim 2^R$ even in the large-*D* limit, showing that correlation length of random tensor networks is not necessarily short. Blue dots are tensors with Gaussian-random entries. Internal bonds have dimension *D*, physical bonds (sticking up at the top) have dimension *d*. The holographic "scale" dimension has size *R*. The correlation length is $\xi \sim 2^R$. The 2*d* version of this network can be contracted by a boundary MPS (dashed box) with the effective bond dimension is $D_{\mathrm{eff}} \sim D^R$, which scales polynomially with ξ.

networks reflect all the important structure found in physical states).

2. Random tensors networks with large correlation lengths

In this section, we construct an ensemble of random 2*d* tensor networks that

(1) have arbitrarily long correlation length ξ,

(2) can be viewed as PEPS with effective bond dimension $D_{\mathrm{eff}} \sim \mathrm{poly}(\xi)$, and

(3) can be reliably analyzed by stat-mech mapping in a large-*D* limit, which predicts that their physical properties can be efficiently computed via an area-law bMPS.

Specifically, inspired by expectation that an RG coarse-graining can reduce a PEPS with any finite correlation length ξ to one with $\xi \lesssim 1$, we consider a shallow generalized multiscale entanglement renormalization ansatz (gMERA) architecture introduced in [48]. A 1*d* version of this structure is shown in Fig. 19, with obvious generalizations to higher-*d*. It consists of a depth, *R*, layers of tensors, in which at layer $1 \leqslant j \leqslant R$, the tensors are connected only at distance 2^j.[4] This shallow gMERA geometry allows one to neatly interpolate between finitely correlated states (*R*-finite) and critical states ($R \rightarrow \infty$). Heuristically, we can view this as a discrete version of the holographic AdS/CFT correspondence, where physical legs live only at the boundary of a (short) extra "scale" dimension, which runs from a short-distance (UV) at the physical boundary, to a longer distance (IR). Here, we con-

sider the case where each tensor in this network has Gaussian random entries, all internal legs have bond-dimension *D*, and physical legs have dimension *d*. Adapting the discussion of typical two-point correlation functions above to this network, one again concludes that the correlations decay exponentially with the size of the smallest domain that includes both sites 1,2. In this network, the smallest domain will run along the IR edge of the network, resulting in

$$C_{\mathrm{typ}}(r) \sim \begin{cases} 1/r^p \mathrm{w}/p \sim \log_2 D; & r < 2^R \\ e^{-r/\xi} \mathrm{w}/\xi \sim 2^R; & r \gtrsim 2^R, \end{cases} \quad (\mathrm{D1})$$

where $\xi \sim 2^R$. We note that, the functional forms listed merely reflect an overall asymptotic decay of the envelop of correlations. In addition, there is a complicated fractal/self-similar modulation inherited from the geometry of the network.

From this expression, we see that the correlation length, ξ can be made as large as desired by controlling the depth of the shallow gMERA. At the same time, this shallow gMERA can be viewed as a PEPS with effective bond dimension $D_{\mathrm{eff}} = D^R = \xi^{\log_2 D}$, which scales polynomially with the correlation length.

The stat-mech mapping of the bMPS entanglement for these shallow gMERA proceeds similarly to that for the 2*d* square PEPS, except that the physical legs arise only at the UV layers. Therefore, for contracting networks representing norms and correlations, the replica-symmetry breaking *e* field is only present in the UV. Nevertheless, this is still sufficient to explicitly break the replica symmetry, and give an area law for the bMPS for any D_{eff}, i.e., for any correlation length ξ.

The area-law coefficient may be estimated in the large-*D* expansion as outlined above for the square-lattice PEPS. The leading contribution again comes form two-site domains of flipped replica spins ($g \neq e$) that straddle the entanglement cut. A new feature is that the straddling domain can occur at any layer in the (shallow) holographic dimension without effecting its free-energy cost, giving rise to an entropic factor that scales as $\sim R$. The resulting bMPS entanglement in the large-*D*, *R* limit is then:

$$S \sim \frac{1}{D^4} R \sim \frac{1}{D^4} \ln \xi. \quad (\mathrm{D2})$$

We note that the physical dimension, *d* does not appear in this expression because the bulk tensors have only virtual legs. At large ξ, the bMPS entanglement in (D2) scales like that of a nearly critical 1*d* system with effective "central charge" $\sim 1/D^4$, with critical scaling cut off by a finite correlation length ξ. In particular,

This example shows that short correlation length is not an intrinsic limitation of random large-*D* tensor networks. However, it remains unclear whether this example fully captures the structure relevant to physical ground states.

[1] S. R. White, Density matrix formulation for quantum renormalization groups, Phys. Rev. Lett. **69**, 2863 (1992).

[2] M. Fannes, B. Nachtergaele, and R. F. Werner, Finitely correlated states on quantum spin chains, Commun. Math. Phys. **144**, 443 (1992).

---

[3] F. Verstraete, V. Murg, and J. I. Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, Adv. Phys. **57**, 143 (2008).

[4] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, Ann. Phys. **326**, 96 (2011).

[5] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, arXiv:cond-mat/0407066.

[6] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational complexity of projected entangled pair states, Phys. Rev. Lett. **98**, 140506 (2007).

[7] J. Haferkamp, D. Hangleiter, J. Eisert, and M. Gluza, Contracting projected entangled pair states is average-case hard, Phys. Rev. Res. **2**, 013010 (2020).

[8] P. Corboz and F. Mila, Crystals of bound states in the magnetization plateaus of the Shastry-Sutherland model, Phys. Rev. Lett. **112**, 147203 (2014).

[9] P. Corboz, T. M. Rice, and M. Troyer, Competing states in the *t-J* model: Uniform *d*-wave state versus stripe state, Phys. Rev. Lett. **113**, 046402 (2014).

[10] I. Niesen and P. Corboz, A tensor network study of the complete ground state phase diagram of the spin-1 bilinear-biquadratic Heisenberg model on the square lattice, SciPost Phys. **3**, 030 (2017).

[11] B.-X. Zheng, C.-M. Chung, P. Corboz, G. Ehlers, M.-P. Qin, R. M. Noack, H. Shi, S. R. White, S. Zhang, and G. K.-L. Chan, Stripe order in the underdoped region of the two-dimensional Hubbard model, Science **358**, 1155 (2017).

[12] B. Ponsioen, S. S. Chung, and P. Corboz, Period 4 stripe in the extended two-dimensional Hubbard model, Phys. Rev. B **100**, 195141 (2019).

[13] J.-Y. Chen, S. Capponi, A. Wietek, M. Mambrini, N. Schuch, and D. Poilblanc, SU(3)$_1$ chiral spin liquid on the square lattice: A view from symmetric projected entangled pair states, Phys. Rev. Lett. **125**, 017201 (2020).

[14] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, Classical simulation of infinite-size quantum lattice systems in two spatial dimensions, Phys. Rev. Lett. **101**, 250602 (2008).

[15] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu *et al.*, Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry, Nat. Commun. **14**, 1952 (2023).

[16] H. Ma, M. Govoni, and G. Galli, Quantum simulations of materials on near-term quantum computers, npj Comput. Mater. **6**, 85 (2020).

[17] T. Guhr, A. Müller-Groeling, and H. A. Weidenmüller, Random-matrix theories in quantum physics: common concepts, Phys. Rep. **299**, 189 (1998).

[18] R. Vasseur, A. C. Potter, Y.-Z. You, and A. W. W. Ludwig, Entanglement transitions from holographic random tensor networks, Phys. Rev. B **100**, 134203 (2019).

[19] R. Levy and B. K. Clark, Entanglement entropy transitions with random tensor networks, arXiv:2108.02225.

[20] Z.-C. Yang, Y. Li, M. P. A. Fisher, and X. Chen, Entanglement phase transitions in random stabilizer tensor networks, Phys. Rev. B **105**, 104306 (2022).

[21] C. Lancien and D. Pérez-García, Correlation length in random MPS and PEPS, Ann. Henri Poincaré **23**, 141 (2021).

[22] R. J. Baxter, Corner transfer matrices in statistical mechanics, J. Phys. A: Math. Theor. **40**, 12577 (2007).

[23] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Unifying projected entangled pair state contractions, New J. Phys. **16**, 033014 (2014).

[24] K. Noh, L. Jiang, and B. Fefferman, Efficient classical simulation of noisy random quantum circuits in one dimension, Quantum **4**, 318 (2020).

[25] Z. Li, S. Sang, and T. H. Hsieh, Entanglement dynamics of noisy random circuits, Phys. Rev. B **107**, 014307 (2023).

[26] P. Hayden, S. Nezami, X.-L. Qi, N. Thomas, M. Walter, and Z. Yang, Holographic duality from random tensor networks, J. High Energy Phys. 11 (2016) 009.

[27] J. Lopez-Piqueres, B. Ware, and R. Vasseur, Mean-field entanglement transitions in random tree tensor networks, Phys. Rev. B **102**, 064202 (2020).

[28] A. Nahum, S. Roy, B. Skinner, and J. Ruhman, Measurement and entanglement phase transitions in all-to-all quantum circuits, on quantum trees, and in landau-ginsburg theory, PRX Quantum **2**, 010352 (2021).

[29] A. Nahum, J. Ruhman, S. Vijay, and J. Haah, Quantum entanglement growth under random unitary dynamics, Phys. Rev. X **7**, 031016 (2017).

[30] A. Nahum, S. Vijay, and J. Haah, Operator spreading in random unitary circuits, Phys. Rev. X **8**, 021014 (2018).

[31] T. Zhou and A. Nahum, Emergent statistical mechanics of entanglement in random unitary circuits, Phys. Rev. B **99**, 174205 (2019).

[32] B. Skinner, J. Ruhman, and A. Nahum, Measurement-induced phase transitions in the dynamics of entanglement, Phys. Rev. X **9**, 031009 (2019).

[33] Y. Li, X. Chen, and M. P. A. Fisher, Quantum Zeno effect and the many-body entanglement transition, Phys. Rev. B **98**, 205136 (2018).

[34] C.-M. Jian, Y.-Z. You, R. Vasseur, and A. W. W. Ludwig, Measurement-induced criticality in random quantum circuits, Phys. Rev. B **101**, 104302 (2020).

[35] Y. Bao, S. Choi, and E. Altman, Theory of the phase transition in random unitary circuits with measurements, Phys. Rev. B **101**, 104301 (2020).

[36] A. C. Potter and R. Vasseur, Entanglement dynamics in hybrid quantum circuits, in *Entanglement in Spin Chains: From Theory to Quantum Technology Applications*, edited by A. Bayat, S. Bose, and H. Johannesson (Springer International Publishing, Cham, 2022), pp. 211–249.

[37] M. P. A. Fisher, V. Khemani, A. Nahum, and S. Vijay, Random quantum circuits, Annu. Rev. Condens. Matter Phys. **14**, 335 (2023).

[38] J. C. Napp, R. L. La Placa, A. M. Dalzell, F. G. S. L. Brandão, and A. W. Harrow, Efficient classical simulation of random shallow 2d quantum circuits, Phys. Rev. X **12**, 021021 (2022).

[39] J. M. Martinis *et al.*, Quantum supremacy using a programmable superconducting processor, Nature (London) **574**, 505 (2019).

[40] T. Nishino and K. Okunishi, Corner transfer matrix algorithm for classical renormalization group, J. Phys. Soc. Jpn. **66**, 3040 (1997).

[41] M. Levin and C. P. Nave, Tensor renormalization group approach to two-dimensional classical lattice models, Phys. Rev. Lett. **99**, 120601 (2007).

[42] R. Orús and G. Vidal, Simulation of two-dimensional quantum systems on an infinite lattice revisited: Corner transfer matrix for tensor contraction, Phys. Rev. B **80**, 094403 (2009).

[43] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, Phys. Rev. Lett. **91**, 147902 (2003).

[44] R. Orús and G. Vidal, Infinite time-evolving block decimation algorithm beyond unitary evolution, Phys. Rev. B **78**, 155117 (2008).

[45] Y. Li, R. Vasseur, M. P. A. Fisher, and A. W. W. Ludwig, Statistical mechanics model for Clifford random tensor networks and monitored quantum circuits, Phys. Rev. B **109**, 174307 (2024).

[46] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, Phys. Rev. A **70**, 052328 (2004).

[47] D. Fattal, T. S. Cubitt, Y. Yamamoto, S. Bravyi, and I. L. Chuang, Entanglement in the stabilizer formalism, arXiv:quant-ph/0406168.

[48] S. Anand, J. Hauschild, Y. Zhang, A. C. Potter, and M. P. Zaletel, Holographic quantum simulation of entanglement renormalization circuits, PRX Quantum **4**, 030334 (2023).