

Size-Driven Quantum Phase Transitions

Johannes Bausch,¹ Toby S. Cubitt,¹ Angelo Lucia,^{2,3,4} David Perez-Garcia,^{2,5,6} and Michael M. Wolf⁷

¹DAMTP, University of Cambridge, Centre for Mathematical Sciences, Cambridge CB3 0WA, United Kingdom

²Departamento de Análisis Matemático, Universidad Complutense de Madrid, 28040 Madrid, Spain

³QMATH, Department of Mathematical Sciences, University of Copenhagen, Universitetsparken 5, 2100 Copenhagen, Denmark

⁴NBIA, Niels Bohr Institute, University of Copenhagen, Blegdamsvej 17, 2100 Copenhagen, Denmark

⁵IMI, Universidad Complutense de Madrid, 28040 Madrid, Spain

⁶ICMAT, Calle Nicolás Cabrera, Campus de Cantoblanco, 28049 Madrid

⁷Department of Mathematics, Technische Universität München, 85748 Garching, Germany

Can the properties of the thermodynamic limit of a many-body quantum system be extrapolated by analysing a sequence of finite-size cases? We present a model for which such an approach gives completely misleading results: a translationally invariant, local Hamiltonian on a square lattice with open boundary conditions and constant spectral gap, which has a classical product ground state for all system sizes smaller than a particular threshold size, but a ground state with topological degeneracy for all system sizes larger than this threshold. Starting from a minimal case with spins of dimension 6 and threshold lattice size 15×15 , we show that the latter grows faster than any computable function with increasing local spin dimension. The resulting effect may be viewed as a new type of quantum phase transition that is driven by the size of the system rather than by an external field or coupling strength. We prove that the construction is thermally robust, opening the possibility that these effects are accessible to experimental observation.

The thermodynamic limit of many-body quantum Hamiltonians is the predominant mathematical tool used to study macroscopic properties of physical systems. A common approach is to analyse a growing sequence of finite system sizes—numerically or experimentally—and then extrapolate the characteristics of interest to the macroscopic limit [38]. This approach has been proven highly successful in numerous cases [2, 19, 23, 33, 39]. On the other hand, it has been shown that e.g. determining whether a system is gapped or gapless in the thermodynamic limit is an undecidable problem [8]. In order to correctly extrapolate the thermodynamic properties of a physical model, it is important to distinguish and recognise features that are a consequence of *finite-size effects*, i.e. properties of the model which are not present in the thermodynamic limit but appear as a by-product of conditions which only hold for systems sizes smaller than some threshold. While some finite-size effects only produce small perturbations of the real model, this is not always the case. For example, relevant finite size effects for the distinct behaviour of antiferromagnets on even or odd system sizes have been proposed in [22] and recently observed experimentally in [12].

In this work we show that finite-size effects can in fact be dominant at arbitrary length scales, to the point of completely obscuring the physics of the thermodynamic limit. This phenomenon occurs not just in pathological examples, but even e.g. for translationally invariant Hamiltonians on low-dimensional spins arranged on a square lattice.

Main result 1: *We explicitly construct models exhibiting the following exotic finite-size effects: below a threshold lattice size with sides of length N , the ground*

state of the Hamiltonian is a non-degenerate product state in the canonical basis, i.e. entirely classical, with a constant spectral gap above it. For system sizes greater than N , however, the low energy space is that of the Toric Code, which is in a sense as quantum as possible: the ground state exhibits topological degeneracy, and the system has anyonic excitations.

Moreover, we provide lower bounds on the threshold lattice size N for spin dimensions ≤ 10 (cf. table I). Already for dimension 10 the threshold size can be as large as $5.2 \cdot 10^{36534}$.

Since in practice in a real-world experiment the ground state cannot be accessed, and only the Gibbs state at some small but non-zero temperature can be prepared, we also prove for one of our models that:

Main result 2: *There exists a finite temperature below which measurement in system sizes smaller than the threshold still yields classical results up to small errors, while the thermodynamic limit converges for low temperatures to the ground state of the Toric Code [7]. Even under a strong fidelity requirement of 10^{-6} , the necessary temperatures are rather mild (cf. table I).*

This sudden and dramatic change in the nature of the ground state may be viewed as a type of *quantum phase transition*, driven by the system size rather than a varying external field or coupling strength.

It has been known for some time that the critical values of external parameters (e.g. temperature, pressure) can depend on the size of the studied samples. Well-studied effects include rising melting points for small particles [6, 10], structural temperature- or pressure-dependent phase

d	4	6	7	8	9	10
N_d	2	15	84	420	$3.3 \cdot 10^7$	$5.2 \cdot 10^{36534}$
$T_d[\frac{\Delta}{k_B}]$	0.058	0.050	0.043	0.038	0.020	5.9μ

TABLE I. Lower bounds on the maximum threshold lattice size $N_d \times N_d$ for different spin dimensions d , after which finite-size effect suddenly disappear and the physics of the thermodynamic limit becomes accessible. Up to dimension 8, a prime periodic Wang tiling gives the lower bound; for larger dimensions, an embedding of Busy Beaver Turing machines. The critical temperature T_d gives an estimate for the temperature at which the transition can still be discriminated with a fidelity of $1 - 10^{-6}$, as a function of the system's spectral gap Δ , which here is equal to the interaction strength since the Hamiltonians are commuting.

transitions between different crystal lattices in thin-film samples and in nano-crystals [20, 25, 35, 40], where the energetically favourable structure differs from that in the thermodynamic limit. And charge density wave order transitions or superconductivity [43, 45], for which the critical temperature changes when approaching mono-layer sample sizes.

Crucially, these works all describe phase transitions driven by an external parameter, whose critical value *varies* depending on the size of the system. Here, we exhibit a transition which is driven by the system size itself; the transition occurs at some critical system size, without any external parameters varying at all. The effects which are most reminiscent of what we prove rigorously here are certain peculiar phenomena for mono-layer samples, or samples with 3 or 13 atom layers, for which the described phase transition cannot be observed anymore [43, 45]; one suggested explanation is a lack of space for nucleation sites [20, 40].

Table I shows an overview of the explicit examples we construct. The threshold system sizes N_d from these examples show that large thresholds are achievable with d -dimensional spins. These are of course lower bounds on the maximum possible threshold size for given local spin dimension; even larger size thresholds may be achievable by other constructions. We chose for concreteness to construct a size-driven transition from classical to Toric Code; our constructions can readily be generalised to instead produce a size-driven transition to other quantum phases.

In order to prove these effects mathematically rigorously, we deliberately construct examples for which there exists an analytic solution. However, this is not true for the general case: as the structure of the Hamiltonian becomes more complex, one expects the behaviour to become more erratic. Indeed, we know that for extremely complex Hamiltonians with very large local spin dimension the behaviour can even become uncomputable [8].




FIG. 1. Plaquette and star interactions of the two-dimensional Toric Code Hamiltonian \mathbf{H}_{TC} . We assign a spin-1/2 particle to every lattice edge marked with a dot. \mathbf{H}_{TC} is a sum of 4-local interactions, plaquettes and stars, which are products of σ^z and σ^x operators, respectively.

It is important to emphasise that the dramatic finite-size effects exhibited here do *not* depend on any careful tuning of coupling strengths, and even occur for Hamiltonians without obvious separation of energy scales in their coupling constants or the matrix entries of the local interactions. Without this restriction, i.e. allowing interactions of magnitude $\mathcal{O}(1)$ and $\mathcal{O}(1/N^2)$, it is in fact trivial to construct a model whose ground state changes character at system size $\mathcal{O}(N)$, with the spectral gap closing as $\mathcal{O}(1/N)$. Our result is much stronger, in the sense that it does not allow such a prediction based solely on an analysis of the coupling strengths, nor from extrapolation of spectral data; in particular, the spectral gap of our model remains constant all the way up to the transition.

Hamiltonian Construction

For local spin dimension $d > 3$, we construct a local, translationally invariant spin Hamiltonian $\mathbf{H}^{(d)}$ on a 2D square lattice with open boundary conditions, such that there exists a threshold system size $N_d \times N_d$, up to which the ground state of $\mathbf{H}^{(d)}$ is entirely classical (i.e. product in the canonical basis), whereas for larger lattice sizes the ground state is that of the Toric Code. Lower bounds on the maximum possible such *transition threshold* N_d for a given local dimension d are shown in table I. For $d > 6$, we give a general procedure for constructing models for which N_d grows faster than any computable function.

The *Toric Code*—introduced by Kitaev [15]—is defined by a Hamiltonian on a two-dimensional spin-1/2 lattice. It is one of the simplest models exhibiting topological order [32, 42].

We start out with a finite lattice as shown in fig. 1. To every edge marked with a dot, we assign a d -dimensional spin $\mathbb{C}^d = \mathcal{H}_{\text{TC}} \oplus \mathcal{H}_{\text{CL}}$ where $\mathcal{H}_{\text{TC}} = \mathbb{C}^2$ and $\mathcal{H}_{\text{CL}} = \mathbb{C}^{d-2}$, such that the overall Hilbert space on the lattice is a tensor product over all separate spins, i.e.

$$\mathcal{H}^{(d)} = \bigotimes (\mathcal{H}_{\text{TC}} \oplus \mathcal{H}_{\text{CL}}) \cong (\bigotimes \mathcal{H}_{\text{TC}}) \oplus (\bigotimes \mathcal{H}_{\text{CL}}) \oplus \mathcal{H}',$$

where \mathcal{H}' contains all mixed \mathcal{H}_{TC} and \mathcal{H}_{CL} terms.

We define a purely classical Hamiltonian \mathbf{H}_{CL} with support only on the subspace $\bigotimes \mathcal{H}_{\text{CL}}$, such that the ground state energy of \mathbf{H}_{CL} is -1 for lattice sizes $N \leq N_d$, and otherwise $\lambda_{\min}(\mathbf{H}_{\text{CL}}) \geq 1/2$. We then combine \mathbf{H}_{CL} with \mathbf{H}_{TC} in such a way that the spectrum below some energy $\lambda' > 0$ is uniquely determined by one or other of these Hamiltonians, by giving an energy penalty for any state with support on \mathcal{H}' . We define the overall Hamiltonian by $\mathbf{H}^{(d)} := \mathbf{H}_{\text{TC}} + \mathbf{H}_{\text{CL}} + \mathbf{H}'$, where

$$\mathbf{H}' := C \sum_{i \sim j} \mathbb{1}_{\text{TC}}^i \otimes \mathbb{1}_{\text{CL}}^j + \mathbb{1}_{\text{CL}}^i \otimes \mathbb{1}_{\text{TC}}^j,$$

where $i \sim j$ stands for a sum over any adjacent spins. $\mathbb{1}_{\text{TC}}$ denotes the projector on the \mathcal{H}_{TC} subspace, and analogously $\mathbb{1}_{\text{CL}}$. Note that \mathbf{H}' only contains 2-local interactions.

In this way, any state $|\psi\rangle \in \mathcal{H}^{(d)}$ supported on \mathcal{H}' will necessarily pick up an energy penalty of at least C . Choosing $C = 1 + \lambda_{\min}(\mathbf{H}_{\text{CL}})$ shifts this part of the spectrum to energies ≥ 1 . We can rescale \mathbf{H}_{TC} to have its low-energy spectrum within $[0, 1/2]$. The ground state of $\mathbf{H}^{(d)}$ will thus be given by either \mathbf{H}_{TC} or \mathbf{H}_{CL} , whichever has the smaller energy. In particular, the system will change abruptly from classical to topologically ordered with anyonic excitations when the lattice size N surpasses the threshold N_d , while keeping a constant spectral gap.

In order to construct a suitable classical Hamiltonian \mathbf{H}_{CL} , we will exploit the same locality structure as in the Toric Code—4-local star and plaquette interactions—since this does not increase the interaction range of the overall Hamiltonian $\mathbf{H}^{(d)}$. We will only consider the case of open boundary conditions, which is the most natural one in this context.

Construction. It is convenient to express the interactions as a so-called *tiling* problem with extra constraints, similar to the well-known Wang tiles. A Wang tile is simply a square tile with coloured edges, and the condition for placing two tiles next to each other is that their edge colours match. Despite this simple setup, it has been shown that the question of whether one can tile the entire plane with a finite set of Wang tiles is in fact undecidable [3], which shows that tiling can encode extremely complex behaviour.

It is easy to represent the tiling problem as a ground state energy problem of a classical, translationally invariant Hamiltonian $\mathbf{H}_{\mathcal{W}}$ on the lattice in fig. 1, and straightforward to verify that this representation only defines a single energy scale. As shown in fig. 3, each tile can be regarded as a plaquette on the lattice. The condition that neighbouring tiles share the same edge colour is thus automatically met. It is clear that for c colours, we need a c -dimensional classical subspace \mathcal{H}_{CL} for each spin, i.e. $d = c + 2$. Working on this classical subspace, we want to find local Hamiltonian plaquette interactions between the spins surrounding a plaquette

p , which we denote with E_p , that penalise any tile not in our set of allowed tiles \mathcal{W} . Specifically, we define a local classical tile interaction via

$$\mathbf{h}^{(p)} := \sum_{w \in \mathcal{W}} a_w \bigotimes_{e \in E_p} |w_{e,p}\rangle \langle w_{e,p}|, \quad (1)$$

where $w_{e,p}$ labels the colour on edge e of tile w placed at plaquette site p . The parameters $(a_w)_{\mathcal{W}}$ do not depend on the plaquette position, and the overall translationally-invariant tiling Hamiltonian is given by the sum over all plaquette sites in the lattice $\mathbf{H}_{\mathcal{W}} := \sum_p (\mathbb{1} - \mathbf{h}^{(p)})$. If $a_w = 1$ for all w , one can show that $\mathbf{H}_{\mathcal{W}}$ has zero energy ground state if and only if the set \mathcal{W} tiles the lattice. If we want to give an energy “bonus” to (i.e. decrease the energy of) a specific tile w , we can set $a_w > 1$. An energy penalty can be given by setting $a_w < 1$. Each tiling thus has a net score—bonuses minus penalties minus mismatching tile pairs. The net score of a specific tiling gives the energy of the corresponding state of $\mathbf{H}_{\mathcal{W}}$. In general, then, the ground state of $\mathbf{H}_{\mathcal{W}}$ will maximise the number of tiles with a bonus while avoiding as many penalties as possible.

A similar construction allows us to add extra star-shaped interactions, constraining tile edges adjacent to a corner. The overall Hamiltonian $\mathbf{H}_{\mathcal{W}} + \mathbf{H}_{\mathcal{S}}$ will then have an optimal ground state in the sense that the sum of penalties minus the sum of bonuses—for both tiles and stars—is minimised. The rigorous argument is presented in lemma 1 in the appendix.

Periodic Tiling. With as few colours c as possible, we create a set of tiles and stars which permit a unique periodic tiling pattern in the ground state. The construction for any number of colours $c \geq 3$ is described in the appendix, and an example for five colours can be seen in fig. 2. One can show that this period grows at least exponentially with the number of colours [26]. By penalising a pattern that occurs precisely once per global period—highlighted in fig. 2—we can ensure that the ground state spectrum jumps to $\lambda_{\min}(\mathbf{H}_{\text{CL}}) \geq 1$ for any larger square size. The transition threshold $N_d = N_{c+2}$ for this model is thus given simply by the horizontal pattern period.

Turing Machine Tiling. Starting from a number of colours $c \geq 6$, it becomes possible to embed a Turing machine into a set of tile and star interactions. We improve on an idea introduced by Robinson [36]—which has been exploited in [8] to show undecidability of the spectral gap—by making use of the extra star constraints to significantly reduce the necessary local dimension. In this new construction, the transition threshold N_d grows faster than any computable function and surpasses the periodic tiling bound for $c \geq 7$.

A *Turing machine* is an abstract machine for algorithmic computation proposed by Turing [41], and is generally accepted as the standard mathematical model for formalising problems of computability and complexity. Such a machine is defined by a finite set of instructions. It is equipped with a finite internal memory, together with a




FIG. 2. Section of the prime periodic pattern for four plus one colours, which is 84-periodic. The left edge and lower corner is enforced by giving the solid black square in the corner a bonus of 1, but penalising black to appear to the right of black on a horizontal edge: this way, the global pattern can be started uniquely with open boundary conditions. The horizontal edge colours form disjoint sets: starting from the bottom row, the colours are red $\{4\}$, blue $\{3\}$, white and yellow $\{1, 2\}$, after which the cycle continues with red. This can be achieved using the 4-local star interactions, e.g. by allowing blue to only appear next to blue and white next to white. For the top row allowing two colours white and yellow, we alternate between them whenever the colour on the vertical edge above it is white. Within each row, these colours on the vertical edge count cyclically through subsequences of length 4, 3 and $4 + 3 = 7$, respectively, which yields the overall horizontal period $\text{lcm}\{4, 3, 7\} = 84$. Every 84 tiles, the pattern necessarily exposes a unique local colour configuration, highlighted in the lower right corner. It can be penalised by a single star interaction and forces the spectrum of the associated Hamiltonian to ≥ 1 when the system size surpasses the threshold $N_5 = 84$.

two-way infinite tape where it can read and write symbols and move the tape left or right. The instructions tell the machine how to update the symbol at the current tape location, and which way to move the tape, depending on the symbol it reads from the tape and its current internal state. In the field of computational complexity, the hardness of a problem is usually defined in terms of the amount of resources needed for a Turing machine to solve it, in terms of time needed and tape consumed.

A Turing machine halts if its internal memory reaches a specified “halting” state, after which no further updates take place. We say a given machine is *halting* if it eventually reaches a halting state. If we restrict to machines with a fixed number of states q for its internal memory, and which read and write only two symbols, i.e. 0 and 1, then the set of possible halting machine programs is finite: there has to exist one that runs for longer, or at least as long as, any other. These machines are called *Busy Beavers*, and their running time is called the *Busy Beaver number* $S(q)$. It is known that $S(q)$ grows faster than any computable function [34] [27].

As in the case of the periodic tiling, we find a way of embedding a Busy Beaver Turing machine into the ground state of a classical Hamiltonian: as soon as the Busy Beaver halts, there will be a penalty, since at that point there is no valid way to continue updating the tape. The tiling is thus possible up to a square size of at least $S(q)/\sqrt{2}$ [28]. As we need $c = q + 2$ colours for a q state Busy Beaver, we immediately find a transition threshold of $N_{q+4} \geq S(q)/\sqrt{2}$.

Thermal Stability

As we show in the appendix, there exists a finite temperature T_d below which the thermal state of the Hamiltonian will still be very close to the ground state for any system size up to the threshold N_d , meaning that any measurement will still reveal a classical state up to very small errors. The temperature T_d depends only inverse-logarithmically on the threshold size, and therefore its scaling with d will be mild in the case of the prime periodic tiling. Table I lists the temperatures corresponding to the various local dimensions d , which is a linear function of Δ/k_B , where Δ is the spectral gap of the Hamiltonian and k_B the Boltzmann constant. Since our models are commuting Hamiltonians, the spectral gap Δ is simply equal to the strength of the interactions between the spins.

For the prime periodic tiling, we also show that if we go to the thermodynamic limit at finite temperature, and then send the temperature to zero (a procedure which is a more mathematically correct description of real implementations [1]), we will recover only ground states of the Toric Code. This shows that there is a complete disagreement between the mathematical predictions from the thermodynamic limit, and any measurement performed on systems below the threshold size. The Busy Beaver construction can be modified in order to show the same property, using the original construction of Robinson [37], at the cost of greatly increasing the local dimension.

Conclusion

By constructing two concrete classes of examples, we have shown that there exist translationally invariant, local Hamiltonians on a 2D square lattice with constant spectral gap and open boundary conditions, which belong to a topologically ordered phase in the thermodynamic limit, but appear to be classical for finite system sizes smaller than a certain threshold. This threshold grows extremely fast as a function of the local spin dimension—for one class it grows faster than any computable function—showing that even for physically realistic systems with low local dimension, erratic behaviour may occur at system sizes that are inaccessible numerically or experimentally. For such systems, physical properties in the thermodynamic limit cannot be extrapolated from sequences of finite-size instances.

The implications of these findings may be profound. Numerical simulations of lattice models play a key role in understanding the dynamics of a system, e.g. in lattice gauge theories [16], fluid dynamics [44] and condensed matter physics [18]. All these simulations are extremely computationally intensive, so accessible lattice sizes are severely limited—e.g. for heavy quark simulations, current lattices have sizes reaching $96^3 \times 192$, the long direction being time [31, Ch. 18]. Our results show that there exist classes of simple, local, physical systems on a lattice of spins with moderate dimension, for which it is impossible to tell with certainty whether the system behaves the same on macroscopic scales as it does on any accessible finite size. In fact, the physical properties of this class of systems will change dramatically above some threshold size, which can even be uncomputable.

Many variations of this problem are possible. It is easy to e.g. reverse our construction and transition from topologically ordered at low system sizes to classical for large lattices, and it is clear that similar constructions using a different Hamiltonian than the Toric Code are possible. As usual when exotic models are found, we expect that the ability to switch properties of a Hamiltonian on and off depending on the system size could also lead to interesting applications in future.

J. B. acknowledges support from the German National Academic Foundation and the EPSRC (grant 1600123). T. S. C. is supported by the Royal Society. A. L. and D. P. G. acknowledge support from MINECO (grant MTM2014-54240-P) and Comunidad de Madrid (grant QUITEMAD+-CM, ref. S2013/ICE-2801). A. L. is supported by MINECO fellowship FPI BES-2012-052404. D. P. G. and M. M. W. acknowledge support the European CHIST-ERA project CQC (funded partially by MINECO grant PRI-PIMCHI-2011-1071). This work was made possible through the support of grant #48322 from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily

reflect the views of the John Templeton Foundation. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 648913).

-
- [1] Michael Aizenman and Elliott H. Lieb. The third law of thermodynamics and the degeneracy of the ground state for lattice systems. *Journal of Statistical Physics*, 24(1):279–297, jan 1981.
 - [2] M. N. Barber. Finite-size scaling. In Cyril Domb and J. Lebowitz, editors, *Phase transitions and critical phenomena Vol. 8*. Academic Press, New York, 1983.
 - [3] R. Berger. *The Undecidability of the Domino Problem*. American Mathematical Society memoirs. American Mathematical Society, 1966.
 - [4] Allen H Brady. The determination of the value of rado's noncomputable function $\sigma(k)$ for four-state turing machines. *Mathematics of Computation*, 40(162):647–665, 1983.
 - [5] Ola Bratteli and Derek W. Robinson. *Operator Algebras and Quantum Statistical Mechanics 2*. Springer Science + Business Media, 1997.
 - [6] Ph. Buffat and J-P. Borel. Size effect on the melting temperature of gold particles. *Physical Review A*, 13(6):2287–2298, jun 1976.
 - [7] Matthew Cha, Pieter Naaijken, and Bruno Nachtergaele. The complete set of infinite volume ground states for kitaev's abelian quantum double models, Aug 2016.
 - [8] Toby S. Cubitt, David Perez-Garcia, and Michael M. Wolf. Undecidability of the spectral gap. *Nature*, 528(7581):207–211, Dec 2015.
 - [9] John Dengis, Robert König, and Fernando Pastawski. An optimal dissipative encoder for the toric code. *New Journal of Physics*, 16(1):013023, jan 2014.
 - [10] A. N. Goldstein, C. M. Echer, and A. P. Alivisatos. Melting in semiconductor nanocrystals. *Science*, 256(5062):1425–1427, jun 1992.
 - [11] Daniel Gottesman and Sandy Irani. The quantum and classical complexity of translationally invariant tiling and hamiltonian problems. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pages 95–104. IEEE, 2009.
 - [12] T. Guidi, B. Gillon, S. A. Mason, E. Garlatti, S. Carretta, P. Santini, A. Stunault, R. Caciuffo, J. van Slageren, B. Klemke, A. Cousson, G.A. Timco, and R.E.P. Winpenny. Direct observation of finite size effects in chains of antiferromagnetically coupled spins. *Nature Communications*, 6:7061, 2015.
 - [13] R. Haag, N. M. Hugenholtz, and M. Winnink. On the equilibrium states in quantum statistical mechanics. *Communications in Mathematical Physics*, 5(3):215–236, jun 1967.
 - [14] M. B. Hastings. Entropy and entanglement in quantum ground states. *Phys. Rev. B*, 76(3), jul 2007.
 - [15] A Yu Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, 2003.
 - [16] John B. Kogut. The lattice gauge theory approach to quantum chromodynamics. *Rev. Mod. Phys.*, 55:775–836, Jul 1983.

- [17] Pavel Kropitz. 6-state 2-symbol #b. http://www.drb.insel.de/~heiner/BB/simKro62_b.html. accessed: April 27th, 2015.
- [18] D.P. Landau, S.P. Lewis, and H.B. Schüttler. *Computer Simulation Studies in Condensed-Matter Physics XIII: Proceedings of the Thirteenth Workshop Athens, Ga, Usa, February 21-25, 2000*. Computer Simulation Studies in Condensed-matter Physics. Springer, 2001.
- [19] L D Landau and E.M. Lifshitz. *Statistical Physics, Third Edition, Part 1: Volume 5 (Course of Theoretical Physics, Volume 5)*. Butterworth-Heinemann, 1980.
- [20] Dehui Li, Gongming Wang, Hung-Chieh Cheng, Chih-Yen Chen, Hao Wu, Yuan Liu, Yu Huang, and Xiangfeng Duan. Size-dependent phase transition in methylammonium lead iodide perovskite microplate crystals. *Nature Communications*, 7:11330, apr 2016.
- [21] Shen Lin and Tibor Rado. Computer studies of turing machine problems. *Journal of the ACM (JACM)*, 12(2):196–212, 1965.
- [22] S. Lounis, Dederichs, P., and S. Blügel. Magnetism of nanowires driven by novel even-odd effects. *Phys. Rev. Lett.*, 101:107204, 2008.
- [23] Norman H. March. *Electron Density Theory of Atoms and Molecules*. Academic Press, London, 1992.
- [24] Heiner Marxen et al. Attacking the busy beaver 5. In *Bull EATCS*, 1990. <http://www.drb.insel.de/~heiner/BB/mabu90.html>, accessed: April 27th, 2015.
- [25] J. M. McHale. Surface energies and thermodynamic phase stability in nanocrystalline aluminas. *Science*, 277(5327):788–791, aug 1997.
- [26] One can show that $N_{q+3} = p_q\# = \Omega(\exp((1 + o(1))q \ln q))$, where $\#$ denotes the primordial function and p_q the q^{th} prime.
- [27] A related quantity, which is also called the Busy Beaver number but is usually denoted $\Sigma(q)$, is defined as the largest number of non-blank symbols written out by the machine before terminating, and is a lower bound on $S(q)$. $\Sigma(q)$ also grows faster than any computable function.
- [28] The constant factor of $\sqrt{2}$ is due to the fact that in our construction the tape of the Turing machine is encoded diagonally with respect to the square lattice, and the head of the machine follows a zig-zag pattern, which in the worst case scenario can only reach a distance of $S(q)/\sqrt{2}$ from the origin.
- [29] Such definition is equivalent to defining a special halting state h for which no further transition is defined.
- [30] One possible way to implement this procedure is to follow a sequence of local steps along an oriented tree structure, as described in [9].
- [31] K.A. Olive et al. Review of Particle Physics. *Chin.Phys.*, C38:090001, 2014.
- [32] Jiannis K Pachos. *Introduction to topological quantum computation*. Cambridge University Press, 2012.
- [33] B. Pirvu, G. Vidal, F. Verstraete, and L. Tagliacozzo. Matrix product states for critical spin chains: Finite-size versus finite-entanglement scaling. *Physical Review B*, 86(7), aug 2012.
- [34] Tibor Rado. On non-computable functions. *Bell System Technical Journal*, 41(3):877–884, 1962.
- [35] Jessie B. Rivest, Lam-Kiu Fong, Prashant K. Jain, Michael F. Toney, and A. Paul Alivisatos. Size dependence of a temperature-induced solid–solid phase transition in copper(i) sulfide. *The Journal of Physical Chemistry Letters*, 2(19):2402–2406, oct 2011.
- [36] Raphael M Robinson. Undecidability and nonperiodicity for tilings of the plane. *Inventiones mathematicae*, 12(3):177–209, 1971.
- [37] Raphael M. Robinson. Undecidability and nonperiodicity for tilings of the plane. *Inventiones Mathematicae*, 12(3):177–209, sep 1971.
- [38] L. Šamaj and Z. Bajnok. *Introduction to the Statistical Physics of Integrable Many-body Systems*. Cambridge University Press, 2013.
- [39] L. Tagliacozzo, Thiago. R. de Oliveira, S. Iblisdir, and J. I. Latorre. Scaling of entanglement support for matrix product states. *Physical Review B*, 78(2), jul 2008.
- [40] S. H. Tolbert and A. P. Alivisatos. Size dependence of a first order solid-solid phase transition: The wurtzite to rock salt transformation in CdSe nanocrystals. *Science*, 265(5170):373–376, jul 1994.
- [41] A. M. Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1):230–265, jan 1937.
- [42] Xiao-Gang Wen. Topological order: From long-range entangled quantum matter to a unified origin of light and electrons. *ISRN Condensed Matter Physics*, 2013:1–20, 2013.
- [43] Xiaoxiang Xi, Liang Zhao, Zefang Wang, Helmuth Berger, László Forró, Jie Shan, and Kin Fai Mak. Strongly enhanced charge-density-wave order in monolayer NbSe₂. *Nature Nanotechnology*, 10(9):765–769, jul 2015.
- [44] Jeffrey Yepez. Quantum lattice-gas model for computational fluid dynamics. *Phys. Rev. E*, 63:046702, Mar 2001.
- [45] Yijun Yu, Fangyuan Yang, Xiu Fang Lu, Ya Jun Yan, Yong-Heum Cho, Liguang Ma, Xiaohai Niu, Sejoong Kim, Young-Woo Son, Donglai Feng, Shiyan Li, Sang-Wook Cheong, Xian Hui Chen, and Yuanbo Zhang. Gate-tunable phase transitions in thin flakes of 1t-TaS₂. *Nature Nanotechnology*, 10(3):270–276, jan 2015.

Embedding a Generalised Tiling into a Hamiltonian Spectrum

In this section, we rigorously formulate the embedding of the tiling problems we consider in this work into the spectrum of a local Hamiltonian. Instead of focusing only on star and plaquette interactions, we take an abstract point of view and define the notion of a *generalised tiling*. Assume $\mathcal{G} = (V, E)$ is a finite undirected graph with coloured vertices, where we allow colours $C := \{1, \dots, c\}$, $c \in \mathbb{N}$. Let $\mathcal{L} := \{l : l \subset \mathcal{G}\}$ be a finite set of (local) interactions, e.g. all the 3- or 4-local star and plaquette interactions on a lattice as in fig. 1. For all interactions $l \in \mathcal{L}$, we allow a finite set of *pieces* $\mathcal{T}_l := \{(c_v)_{v \in l}\}$ —where the family $(c_v)_{v \in l}$ assigns a colour to every vertex in l —and a weight function $w_l : \mathcal{T}_l \rightarrow \mathbb{R}$. Now assign a colour to each vertex in \mathcal{G} , e.g. by defining a family $(c_v)_{v \in V}$, $c_v \in C$. The *score* of this assignment is then given by

$$\text{score} := \sum_{l \in \mathcal{L}} \begin{cases} 1 - w_l(t_l) & \text{if } (c_v)_{v \in l} \text{ is a valid piece in } \mathcal{T}_l \\ 1 & \text{otherwise.} \end{cases}$$

For $w_l(t_l) < 1$, we can thus give a score penalty, and $w_l(t_l) > 1$ gives a bonus to piece t_l at site l . An assignment $w_l(t_l) = 1$ is neutral and gives neither bonus nor penalty. Observe that *not* including a piece in the piece set \mathcal{T}_l is equivalent to giving it a weight of 0. It is easy to see how this specialises to our tiling examples: in case of the periodic tiling and for l a plaquette interaction in the bulk, the sets \mathcal{T}_l would all be identical and correspond to the allowed 4-local tiles. The w_l then assign the bonuses or penalties, accordingly.

We formulate the following lemma.

Lemma 1. *Define a Hilbert space $\mathcal{H} := \bigotimes_{v \in V} \mathbb{C}^c$ over the interaction graph \mathcal{G} , assigning c -dimensional qudits to each vertex $v \in V$. Then there exists a classical Hamiltonian \mathbf{H} on \mathcal{H} , diagonal in the computational basis, with \mathcal{L} -local interactions such that the eigenvalue λ for a basis state $|\psi\rangle = \bigotimes_{v \in V} |c_v\rangle$ is given by the score of the associated generalised tiling, i.e.*

$$\lambda = \sum_{l \in \mathcal{L}} \begin{cases} 1 - w_l(t_l) & \text{if } |\psi\rangle|_l \in \mathcal{T}_l \\ 1 & \text{otherwise.} \end{cases}$$

We denote with $|\psi\rangle|_l$ the restriction of $|\psi\rangle$ to the subspace $\bigotimes_{v \in l} \mathbb{C}^c \leq \mathcal{H}$.

Proof. Define

$$\mathbf{H} := \sum_{l \in \mathcal{L}} \left(\mathbb{1} - \sum_{t \in \mathcal{T}_l} w_l(t) \Pi_t \right),$$

where $\Pi_t := \bigotimes_{v \in l} |t_v\rangle\langle t_v|$ denotes the projector onto the valid piece $t \in \mathcal{T}_l$ for interaction $l \in \mathcal{L}$, and t_v denotes the

colour of vertex v for piece t . Take a computational basis state $|\psi\rangle = \bigotimes_{v \in V} |c_v\rangle$. Then

$$\begin{aligned} \mathbf{H}|\psi\rangle &= \sum_{l \in \mathcal{L}} \left(|\psi\rangle - \sum_{t \in \mathcal{T}_l} |\psi\rangle \begin{cases} w_l(t_l) & \text{if } |\psi\rangle|_l \in \mathcal{T}_l \\ 0 & \text{otherwise} \end{cases} \right) \\ &= \lambda |\psi\rangle, \end{aligned}$$

and the claim follows. \square

This allows us to conclude the following corollary.

Corollary 2. *The ground state energy of \mathbf{H} is determined by the lowest score assignment of the associated generalised tiling problem.*

Equipped with this machinery, it suffices to formulate generalised tiling problems on the square lattices as in fig. 1 with 3- and 4-local interactions, such that for lattice sizes below some threshold, the lowest score assignment has a score $\leq -1/2$, and above the threshold the lowest score assignment has a score ≥ 1 . This way, combining the Toric Code Hamiltonian \mathbf{H}_{TC} via lemma 3 creates a model with a size-induced transition from classical to topological ground state. Observe that we require our model to be translationally invariant.

The Toric Code

The Toric Code Hamiltonian \mathbf{H}_{TC} is a sum of 3- and 4-local interactions

$$\mathbf{H}_{\text{TC}} := -J \sum_s \mathbf{A}^{(s)} - J \sum_p \mathbf{B}^{(p)},$$

with $\mathbf{A}^{(s)} := \prod_{i \in s} \sigma_i^x$ a product of Pauli σ^x acting on 4 spins i adjacent to vertex s as seen in fig. 1. The $\mathbf{B}^{(p)} := \prod_{i \in p} \sigma_i^z$ are defined analogously. We call the $\mathbf{A}^{(s)}$ *star* and the $\mathbf{B}^{(p)}$ *plaquette*-interactions, respectively. The free parameter $J > 0$ is a coupling strength and can be used to rescale the spectrum.

Prime Period Tiling

The key idea is to create a tiling pattern that can tile the entire plane with a very large period p . We require that a certain locally detectable sub-pattern—i.e. using a star interaction—occurs exactly once per period. By disallowing this sub-pattern, the tiling will be possible up to a square of size $p \times p$, but once the grid surpasses this size, there will be at least one pattern violation, which can be penalised locally with a Hamiltonian term.

General Construction. For the general construction, we first regard the following discrete optimisation problem. Assume we have q colours available. We want to construct a family of tuples $(r_i)_{1 \leq i \leq f}$, each of which stands for a row of colours $r_i = (r_{i1}, r_{i2}, \dots, r_{im_i})$. These rows have to satisfy three constraints.

1. There are fewer than q rows overall, i.e. $f \leq q$.
2. Each row has fewer than q colours, i.e. $m_i \leq q \forall i$.
3. For the first and last row, each colour r_{ij} is picked from the q colours available, i.e. $r_{1j}, r_{fj} \in \{1, \dots, q\}$ —for all other rows, we leave out the last, i.e. $r_{ij} \in \{1, \dots, q-1\}$.

We can associate a period $p_i := \sum_{j=1}^{m_i} r_{ij}$ to each row i . The rows r_i are now chosen such that the objective function $p(q) := \text{lcm}\{p_i : i \leq f\}$ —i.e. the overall period—is maximised.

We now give a description on how to translate such an optimal row family $(r_i)_i$ into a set of tiles and stars that enforces a unique horizontal tile sequence with periodicity $p(q)$. More specifically, for each row, we define tiles that allow a colour pattern

$$1, \dots, r_{i1}, 1, \dots, r_{i2}, \dots, 1, \dots, r_{i(m_i-1)}, 1, \dots, r_{im_i} \quad (2)$$

on their vertical edge—i.e. the tiles form sub-periods, starting at 1 and counting to r_{ij} , then starting at 1 again, counting to $r_{i(j+1)}$ etc.. Cleverly choosing colours on the horizontal edges then a) make this pattern unique for each row, and b) enforce a unique stacking order of the rows overall—which in turn yields a $p(q)$ -periodic global tile pattern.


To facilitate the tile notation, we use a few shorthands.

- The last sub-period for each row has highest colour $r_{im_i} =: \bar{r}_i$.
- We sequentially enumerate the sub-periods with colours for use on the horizontal edges, i.e. $r_{11} \leftrightarrow 1, r_{12} \leftrightarrow 2, \dots, \bar{r}_1 \leftrightarrow m_1, r_{21} \leftrightarrow m_1 + 1$ etc.. The highest such label for every row is denoted with h_i , and the lowest l_i , e.g. for the first row $l_1 = 1$ and $h_1 = m_1$. More rigorously, we have the sequences $h_0 = 0, h_i = h_{i-1} + m_i$ and $l_i = h_{i-1} + 1$.
- The set of colours on the horizontal edges needed for the i^{th} row is denoted with $V_i := \{l_i, \dots, h_i\}$, respectively.

For every row r_i , we then define the tiles


$$\begin{array}{|c|c|c|} \hline & t & \\ \hline c & & c' \\ \hline & b & \\ \hline \end{array} \quad \begin{array}{l} \forall t \in V_{(i-1 \bmod f)}, \quad b \in V_{(i \bmod f)} \\ \forall c = 1, \dots, r_{ij}, \end{array}$$

where $c' = c + 1 \bmod r_{ij}$. As an example, consider the first row with $V_1 = \{1, \dots, m_1\}$. We obtain a set of tiles




where $*_f$ stands for any colour allowed on the bottom of the last row, i.e. $*_f \in V_f$. All other rows are defined analogously, where the top and bottom colours are chosen successively, i.e. for the i^{th} row, we use colours V_i for the bottom and any V_{i-1} for the top.

On their own, the tiles from different sets can be mixed at will. To enforce that each row can only be assembled from its own tile set, for each row i , we restrict to the following star configurations:




where $j' := l_i + (j + 1 \bmod m_i)$. It is easy to verify that each row defines a unique p_i -periodic horizontal tile pattern as in eq. (2). As the top colours of the i^{th} row are restricted to the bottom colours of the $(i-1)^{\text{th}}$ row—modulo the numbers of rows f —the rows can be stacked above each other uniquely. Every block of rows r_1, \dots, r_f , stacked vertically, thus defines a valid horizontally p -periodic tiling pattern for the plane.

In order to be able to detect this periodicity locally, we make use of the extra colour available in all but the first and last row due to constraint (3). For all $i = 2, \dots, f-1$, we add two tiles




Alternative to the row sequence $\dots, \bar{r}_i - 1, \bar{r}_i, 1, 2, \dots$, this allows counting $\dots, \bar{r}_i - 1, \bar{r}_i, \mathbf{q}, 2, \dots$. By adding the star penalties




we ensure that whenever two consecutive rows complete a cycle in the same column, we mark the occurrence with a q instead of a 1. This way, if in the first row we finish a cycle with a 1 and observe a q right below, we know that the entire horizontal pattern has completed one period. To be more specific, every $p(q) = \text{lcm}\{p_1, \dots, p_K, p_L, p_f\}$

tiles, where $L = f - 1$ and $K = f - 2$, we have the pattern




We call this sub-pattern a *period marker* and penalise it with a weight of 2.

So far we have constructed a tile set which can periodically tile the entire plane. By disallowing said period marker, we restrict the tileable square size to at most $p \times p$. Observe however that due to the freedom to shift sets of rows horizontally and the entire pattern vertically, there are a potentially huge number of possibilities to tile any square smaller than $p \times p$. We will thus add a special colour to fix this freedom, borrowing an idea from [11]. We will enforce a specific pattern in the lower left corner, which uniquely fixes the starting configuration for the bulk, without imposing any boundary condition, but instead by adding bulk interactions which will have the effect of favouring the desired configuration in the boundary. We add the following tiles:



We further disallow black appearing to the right of black using a star constraint, and give a bonus of $1/2$ to the all-black tile. It is then easy to verify that the best score tiling starts with the following configuration in the lower left corner:



It is straightforward to verify that starting from this corner configuration, the plane can be tiled uniquely up to a grid size of $p \times p$ with a net score of $-1/2$, after which the net penalty jumps to a value ≥ 1 .

In table II, we list the cases $q = 2, \dots, 8$ with a solution to the associated constraint problem and the resulting overall period p . It is easy to see that in the case of 2 colours only, the extra black tile to remove degeneracies is redundant.

Five Colour Tiling Example. We give the $q = 4$ colour case as an explicit example. The tile set in fig. 3 defines three disjoint sets of tiles, each of which can be assembled into horizontal lines, which in turn can be stacked above

colours $q + 1$	line periods	overall period p
2^\dagger	2, 2	4
4	3, 5	15
5	4, 3, 7	84
6	5, 4, 3, 7	420
7	6, 5, 7, 11	2310
8	7, 6, 5, 11, 13	30030
9	8, 7, 11, 13, 15	120120

TABLE II. Maximum tiling periods for a given number of colours (plus one special black colour needed for $q > 2$). The second column shows how to distribute the periods between the lines. † For 2 colours, it is easy to see that the extra black tile is redundant.





FIG. 3. Four+1 colour tile set that defines a tiling of the plane with period 84. The striped top in the second row denotes either of the colours 1 or 2.

each other in a unique order. To avoid mixing the tiles from different sets on one same line, we add the following star operators with parameter $b_s = 1$




The third row then unambiguously assembles to a line with a horizontal period of 4, while the vertical edges appearing on the first line periodically cycle through 1, 2, 3, 4, 1, 2, 3 with a period of 7. The horizontal edges of the second line are fixed by the line above and below, but there exists some freedom to choose the colours on the vertical edge. More specifically, we use the freedom of either counting $\dots, 2, 3, 1, 2, \dots$ or $\dots, 2, 3, 4, 2, \dots$ to detect when all three lines complete a period in the same

states $ Q $	colours c	$S(Q)$	threshold N_d
3	6	21 [21]	14
4	6	107 [4]	75
5	7	$\geq 4.7 \cdot 10^7$ [24]	$3.3 \cdot 10^7$
6	8	$\geq 7.4 \cdot 10^{36534}$ [17]	$5.2 \cdot 10^{36534}$

TABLE III. Number of 2 symbol Turing machine states Q and tile colours $c = \max\{|A|^2 + 2, |A| + |Q|\}$ required for the embedding. $S(|Q|)$ is the Busy Beaver number, and N_d denotes the lower bound on the maximum lattice threshold size, where $d = c + 2$.

column. We add a penalty for the configuration




by adding the corresponding operator multiplied by $b_s = -1$, which enforces colour 4 to appear instead of colour 1 whenever the first row finishes one cycle at the same time as the second one below. The combined period p of the three lines is thus given by $\text{lcm}(7, 4, 3) = 84$, and it can be detected by penalising the configuration



with a penalty of 2 (i.e., with $b_s = -2$.)

The freedom to horizontally shift the lines relative to each other or the entire pattern vertically is fixed without adding boundary conditions with the special tiles



We choose $a_w = 2$ for the first. Starting from there, the entire plane can be tiled uniquely up to a grid size of 84×84 , after which the penalised star—eq. (3)—naturally occurs and the net penalty is ≥ 1 . A section of the complete 5-colour tiling can be seen in fig. 2.

Generalising the prime tiling to higher dimensions, we obtain the periods as given in table II.



Turing Machine Tiling

A Turing machine is given by finite sets of states Q and symbols A with a transition function $\delta : Q \times A \rightarrow Q \times A \times \{\text{left}, \text{right}\}$, representing the set of instructions of the Turing machine. The machine is equipped with a tape, which is sequence of symbols arranged in a 1-dimensional line extending indefinitely in both directions,


and initialised with a special “blank” symbol (which we will denote 0 for simplicity of notation). The machine has an internal state $q \in Q$ and a head which sits over one of the symbols of the tape: at each step, the head reads the symbols s underneath, and it will write the symbol \bar{s} , change its internal state to \bar{q} and then move in direction $d \in \{\text{left}, \text{right}\}$, where $(\bar{q}, \bar{s}, d) = \delta(q, s)$. The machine starts in an initial state $q_0 \in Q$ and *halts* if there is no forward transition for a given tuple (q, s) [29].

We will show how to construct a set of plaquette and star interactions in such a way that the ground state encodes the history of a run of a given Turing machine. The construction will involve the use of Wang tilings and of star interactions: the latter will allow us to greatly reduce the spin dimension needed by previous works which were based only on the tiling problem [3, 36].

We assume the lattice is oriented as in fig. 4, where the Turing machine starts in the lower left corner. We chose to encode the tape of the Turing machine at a given time step in diagonal direction across the square lattice (denoted by the thick gray lines in fig. 4), and movement in the orthogonal direction represents the time evolution. Moreover, we store the head position and internal state of the Turing machine by including the state q in the tape, on the right of the symbol which will be read by the machine. Using this convention, we have that—even if the tape space is finite—it is extended by one symbol in both directions at each time step, and therefore the tape available will always be sufficient for the machine to run.

We will interpret colours on horizontal and vertical edges differently—horizontal either as pair of symbols $s_1 s_2 \in A \times A$ or special boundary colour  or , and vertical either as symbol $s \in A$ or state $q \in Q$, the latter of which we highlight in red.

As we did in the periodic tiling construction, we use special bulk interaction (which will be present everywhere in the lattice) to constrain the left and bottom boundaries. In order to do so, we use the following interactions



Note that the 2-local blue term is indeed a bulk interaction, and not a “cut-off” boundary term. By giving the last of these tiles a single bonus of 2—similar to the all black tile for the periodic tiling—we obtain precisely one choice for the left and lower grid edges, namely all 0s as in fig. 4; in particular, the last shown tile with initial state q_0 correctly initialises the Turing machine in the lower left corner. This valid initial configuration defines the unique highest-net-bonus tiling possible.

To avoid cases where we validly tile the plane without a TM head—i.e. with net bonus zero—we use a star interaction to give a bonus of $1/2$ for any *white* symbol on a vertical edge appearing to the left or right of another

arbitrary symbol, i.e.

$$1 \times \textcircled{**} - \textcircled{**} \quad , \quad \frac{1}{2} \times \textcircled{**} - \textcircled{00} \quad \text{and} \quad \frac{1}{2} \times \textcircled{00} - \textcircled{**}.$$

We further give a penalty of 1 for the white symbol appearing anywhere. This way, in the bulk, the net contribution of $2 \times 1/2 - 1 = 0$ for each of the white edges, whereas if they appeared on the left end of the plane a net penalty $\geq 1/2$ would be inflicted. A similar combination of bonus and penalty terms allows us to ensure that the lower edge is blue, and all other configurations obtain a net penalty of $\geq 1/2$ as well. Like that, there exists no configuration with net penalty $< 1/2$ without the initial q_0 tile in the lower left corner. From now on, we treat the boundary symbols $\textcircled{00}$ and $\textcircled{0\bar{0}}$ as equivalent to $\textcircled{00}$.

To implement the transitions rules we effectively need 6 different spins to interact (three for each time step): in fact, if the tape around the head reads s, q, t for some $q \in Q$ and some $s, t \in A$, then it has to be updated to \bar{q}, \bar{s}, t if $\delta(q, s) = (\bar{q}, \bar{s}, \text{left})$, while it has to be updated to \bar{s}, t, q if instead $\delta(q, s) = (\bar{q}, \bar{s}, \text{right})$.

Since we only have at our disposition 4-body interactions, in order to implement an effective 6 body interaction we will make use of the extra register we have allowed for in the horizontal edges, which will allow to “synchronise” a plaquette and a star interaction, as shown in fig. 4. This is done by defining, for every transition, a pair of tiles and stars, i.e. if $\delta(q, s) = (\bar{q}, \bar{s}, \text{left})$

$$\text{Star} \quad \text{and} \quad \text{Plaquette} \quad (4)$$

where $*$ represents any symbol, and for an analogous right transition

$$\text{Star} \quad \text{and} \quad \text{Plaquette} \quad (5)$$

Observe how the symbol pairs are necessary to uniquely couple the pair of interactions to obtain the left and right transition depicted in fig. 4, but are disregarded for any successive transition. The rest of the tape which is not affected by the transition rules has to be copied verbatim to the next time step. Implementing such bookkeeping tiles is straightforward, as we only need to take care of the extra—and in this situation unused—register in the horizontal edges, which is discarded when copying to vertical edges, and set to the “blank” symbol when copying from vertical to horizontal edges. More precisely, for every $a, b \in A$, we define

$$\text{Star} \quad \text{and} \quad \text{Plaquette} \quad (6)$$

Overall, this construction thus requires $c = \max\{|A|^2 + 2, |A| + |Q|\}$ colours. It is easy to verify that starting from the initial tile, a square can be uniquely tiled with net bonuses 1 if and only if the Turing machine does *not* halt within its boundaries. All other tilings necessarily violate at least one constraint and thus have a net penalty $\geq 1/2$. A sample evolution can be seen in fig. 4.

The maximum number of steps any *halting* Turing machine with $|Q|$ states and 2 symbols can take before halting is called the *Busy Beaver number* and is denoted by $S(|Q|)$. Defined in [34], it is known to grow faster than any computable function. The staggering threshold sizes N_d in table I show that there is no hope to address the question of extrapolating physical properties of a general system solely with an increase in computational power.

Combining Hamiltonian Spectra

Lemma 3. *Let \mathbf{H}_1 and \mathbf{H}_2 two local Hamiltonian defined on $\bigotimes_{u \in \Lambda} \mathbb{C}^{d_1}$ and $\bigotimes_{u \in \Lambda} \mathbb{C}^{d_2}$ for some interaction graph Λ . Let further $\mu \in \mathbb{R}$. Then there exists a Hamiltonian \mathbf{H} on $\mathcal{H} = \bigotimes_{u \in \Lambda} \mathbb{C}^{d_1} \oplus \mathbb{C}^{d_2}$ with the following properties:*

1. *Any eigenvector v of \mathbf{H} with eigenvalue $\lambda \leq \mu$ is given by an eigenvector of either \mathbf{H}_1 or \mathbf{H}_2 , extended canonically to the larger Hilbert space \mathcal{H} , with the same eigenvalue λ .*
2. *\mathbf{H} is translationally invariant if \mathbf{H}_1 and \mathbf{H}_2 are.*
3. *\mathbf{H} contains nearest neighbour interactions and otherwise leaves the interaction range of \mathbf{H}_1 and \mathbf{H}_2 intact.*

Proof. Let $\mathbb{1}_1$ and $\mathbb{1}_2$ be the identity operators on \mathbb{C}^{d_1} and \mathbb{C}^{d_2} , respectively. Let $\delta := 1 + \mu$. Define further

$$\mathbf{H}_0 := \delta \sum_{i \sim j} \mathbb{1}_1^i \otimes \mathbb{1}_2^j + \mathbb{1}_2^i \otimes \mathbb{1}_1^j,$$

where $i \sim j$ denotes any neighbouring spin pairs. Set $\mathbf{H} := \mathbf{H}_0 + \mathbf{H}'_1 + \mathbf{H}'_2$, where $\mathbf{H}'_1 := \mathbf{H}_1 \oplus \mathbf{0}$ and analogously for \mathbf{H}'_2 .

The last two claims are satisfied by construction. To prove the first point, note that \mathbf{H}_0 , \mathbf{H}'_1 and \mathbf{H}'_2 commute and thus share a common eigenbasis with spectrum $\sigma(\mathbf{H}) = \sigma(\mathbf{H}_0) + \sigma(\mathbf{H}_1) + \sigma(\mathbf{H}_2)$. Since $\delta > \mu$, any eigenstate of \mathbf{H} with eigenvalue $\lambda \leq \mu$ thus has to be in the kernel $\ker \mathbf{H}_0 \equiv \text{supp}(\mathbf{H}'_1 + \mathbf{H}'_2) = \text{supp} \mathbf{H}'_1 \sqcup \text{supp} \mathbf{H}'_2$, and the claim follows. \square

THERMAL STABILITY

Stability up to Transition Threshold Size

We will now show that for both the periodic tiling and the Busy Beaver model there exists a finite in-




FIG. 4. Embedding of a Turing machine into a tiling problem with extra star constraints. We chose a representation in which the Turing machine head sits in between the tape symbols, reading and writing the symbol on its left. Every grey horizontal slice shows the tape at one evolution step: it is initialised to $\dots, 0, q_0, 0, \dots$, where q_0 is the initial machine state, and every successive step is uniquely defined by the transition rules. Shown here is the 2 state Busy Beaver which halts after 6 steps. Since there is no valid tile with a halting state, the system necessarily frustrates for lattices larger $6/\sqrt{2} \approx 4$ tiles on each edge.

Each right transition $(q_i, a) \mapsto (q_f, a', \text{right})$ or left transition $(q_i, a) \mapsto (q_f, a', \text{left})$ is translated into a pair of 4-local plaquette and star interactions, as depicted to the right. Observe how in both cases the tile part of the interaction has to know the initial symbol a , which is why the star operator creates a temporary copy of it. This copy is shown as small symbols and numbers to the right of the actual tape content and ignored in any following transition.

As the available space grows by one symbol in both directions at each step, there is always enough tape available for the Turing machine. The coloured terms are used to initialise this spare tape to 0. Away from the head, additional interactions are used to copy currently unused tape segments forward. The exact construction with all interaction terms is explained in detail in the appendix.

verse temperature β_d (depending on the local dimension due to its explicit dependence on the threshold size N_d), above which the thermal state of the Hamiltonian $\rho_\beta = \exp(-\beta \mathbf{H}^{(d)})/Z_\beta$ will still be very close to a classical state, i.e. to the classical ground state of $\mathbf{H}^{(d)}$.

We will recall the following observation of Hastings [14]: if ρ_0 is the density matrix corresponding to the ground state P_0 of \mathbf{H} (i.e. $P_0 = (\text{tr } P_0)\rho_0$), then we have the following bound:

$$\begin{aligned} \|\rho_\beta - \rho_0\|_1 &= \text{tr} \left| \frac{e^{-\beta \mathbf{H}} - Z_\beta \rho_0}{Z_\beta} \right| = \text{tr} \left| \frac{e^{-\beta \mathbf{H}} - e^{-\beta \lambda_0} P_0}{Z_\beta} + \frac{e^{-\beta \lambda_0} P_0 - Z_\beta \rho_0}{Z_\beta} \right| \\ &\leq \frac{\text{tr} |e^{-\beta \mathbf{H}} - e^{-\beta \lambda_0} P_0|}{Z_\beta} + \frac{\text{tr} |e^{-\beta \lambda_0} P_0 - Z_\beta \rho_0|}{Z_\beta} = 2 \frac{|Z_\beta - e^{-\beta \lambda_0} \text{tr } P_0|}{Z_\beta} \end{aligned}$$

where λ_0 is the ground state energy. Moreover, since $Z_\beta \geq \text{tr } P_0 e^{-\beta \lambda_0} \geq e^{-\beta \lambda_0}$, we have that

$$\frac{|Z_\beta - \text{tr } P_0 e^{-\beta \lambda_0}|}{Z_\beta} \leq \sum_{\lambda \in \sigma(\mathbf{H}) \setminus \lambda_0} e^{-\beta(\lambda - \lambda_0)}, \quad (7)$$

where eigenvalues are counted with their multiplicity.

Let us denote by Δ the spectral gap of \mathbf{H} , and by $\eta(m)$ the number of eigenstates with energy in the range $[m\Delta + \lambda_0, (m+1)\Delta + \lambda_0)$. Then, following [14], if \mathbf{H}

satisfies

$$\eta(m) \leq \frac{K^m}{m!}, \quad (8)$$

then we can bound the r.h.s. of eq. (7) by

$$2 \sum_{m=1}^{\infty} \eta(m) e^{-\beta \Delta m} \leq 2 \sum_{m=1}^{\infty} \frac{(K e^{-\beta \Delta})^m}{m!} = 2(e^{K e^{-\beta \Delta}} - 1).$$

Since in our case $\mathbf{H}^{(d)}$ is a commuting Hamiltonian, $\eta(m)$ grows as $\binom{N_d^2}{m} \leq N_d^{2m}/m!$, which implies

$$\|\rho_\beta - \rho_0\|_1 \leq 2(e^{N_d^2 e^{-\beta \Delta}} - 1). \quad (9)$$

Fix a small $\epsilon > 0$, and let us now choose β_d such that

$$2(e^{N_d^2 e^{-\beta_d \Delta}} - 1) \leq \epsilon,$$

where N_d is the critical system size of $\mathbf{H}^{(d)}$, meaning that

$$\beta_d = \frac{1}{\Delta} [2 \log N_d - \log \log (1 + \frac{\epsilon}{2})]. \quad (10)$$

With this choice of β_d , then we have that for all system sizes $N \leq N_d$ and all $\beta \geq \beta_d$, the thermal state ρ_β is ϵ -close to the ground state of $\mathbf{H}^{(d)}$, which as we have seen is a classical product state.

On the other hand, if $N > N_d$, from the periodic tiling construction we see that the sector of $\mathbf{H}^{(d)}$ corresponding to the tiling Hamiltonian \mathbf{H}_{CL} necessarily picks up a energy penalty every period of N_d (it actually picks up even more, given that the pattern is repeated vertically with a period corresponding to the number of colours, and therefore the energy penalty of every $N_d \times N_d$ square is at least N_d/d). This implies that every eigenstate of \mathbf{H}_{CL} has a strictly positive energy density, and the spectrum of \mathbf{H}_{CL} is contained in $[(N/N_d)^2, \infty)$. This is not true for the Busy Beaver embedding, as it could be more favourable to terminate the computation, and then simply continue with a blank tape. The energy density decreases to zero in this case. Augmenting the construction with a base layer formed from Robinson tiles, it is however possible to make this Turing machine embedding similarly robust. We refer the reader to [37] and [8, ch. 8] for more details.

Thermodynamic Limit

Let us now recall that a state in the thermodynamic limit is given by a linear, positive and normalised functional ω on the algebra of quasi-local observables \mathcal{A} , which is the (norm closure of the) inductive limit of the finite matrix algebras $\mathcal{B}_\Lambda = \mathcal{B}(\otimes_{i \in \Lambda} \mathcal{H}_i^{(d)})$, where Λ is an ascending sequences of finite lattices converging to \mathbb{Z}^2 [1, 5].

Given a local Hamiltonian \mathbf{H} and a finite region Λ , we define its (exterior) boundary $\partial\Lambda$ as the set of sites in the complementary of Λ for which there is an interaction




FIG. 5. Boundary $\partial\Lambda$ of a smooth-bounded rectangular region Λ of the spin lattice. Depicted in red are spins that share a plaquette interaction with at least one spin within Λ , and in blue the ones which also share a star interaction. By definition $\partial\Lambda \cup \Lambda = \emptyset$.

term in \mathbf{H} acting nontrivially on sites of Λ and $\partial\Lambda$ simultaneously, as shown in fig. 5. $\bar{\Lambda}$ is defined as $\Lambda \cup \partial\Lambda$ and, for a region R , \mathbf{H}_R will denote the restriction of \mathbf{H} to all interactions which are totally contained in R . A *ground state* is then defined as a state functional ω , such that for any finite Λ , and any local observable $A \in \mathcal{B}_\Lambda$,

$$\omega(A^\dagger [\mathbf{H}_{\bar{\Lambda}}, A]) \geq 0. \quad (11)$$

This definition can be obtained by taking the zero temperature limit in the definition of finite temperature equilibrium states as defined by the KMS condition (i.e. the limit of increasing-volume Gibbs ensembles satisfying the KMS condition, cf. [13, eq. 4.2]). Loosely speaking, it expresses the intuitive understanding that any local perturbation should not decrease the energy of a ground state (cf. [7]).

Note that since both A and $\mathbf{H}_{\bar{\Lambda}}$ have finite support, it is possible to rewrite eq. (11) in terms of the reduced density matrix of ω over $\bar{\Lambda}$, which we denote by $\rho_{\bar{\Lambda}}$:

$$0 \leq \omega(A^\dagger [\mathbf{H}_{\bar{\Lambda}}, A]) = \text{tr}(\rho_{\bar{\Lambda}} A^\dagger [\mathbf{H}_{\bar{\Lambda}}, A]),$$

or equivalently

$$\text{tr}(\rho_{\bar{\Lambda}} A^\dagger \mathbf{H}_{\bar{\Lambda}} A) \geq \text{tr}(\rho_{\bar{\Lambda}} A^\dagger A \mathbf{H}_{\bar{\Lambda}}) \quad (12)$$

for all Λ and all $A \in \mathcal{B}_\Lambda$. In turn, this implies that

$$\text{tr}(\Phi(\rho_{\bar{\Lambda}}) \mathbf{H}_{\bar{\Lambda}}) \geq \text{tr}(\rho_{\bar{\Lambda}} \mathbf{H}_{\bar{\Lambda}}), \quad (13)$$

for any completely positive, trace preserving linear map Φ , as can be seen by applying eq. (12) to the Kraus operators of $\Phi(\cdot) = \sum_i A_i \cdot A_i^\dagger$.

We will argue that the only ground states of the periodic tiling Hamiltonian $\mathbf{H}^{(d)}$ are the ground states of the Toric Code. This in turns implies, that if we take first the limit of N going to infinity, and then we send the temperature to zero, we recover only ground states of the Toric Code.

Key to our argument is that part of our Hamiltonian—i.e. \mathbf{H}_0 —is a ferromagnetic Ising-type interaction, where

spin up and down are now the tiling and Toric code subspaces, respectively. We will follow the same proof technique used to show that the 2D Ising model with an external magnetic field has a unique ground state [1, ex. 5] to show that any ground state in the thermodynamic limit of our model is completely in the Toric code subspace, by which we mean that for all Λ , $\omega(\Pi_{\text{TC},\Lambda}^\perp) = 0$ for the projector onto the Toric code subspace $\Pi_{\text{TC},\Lambda}$ supported on Λ .

Let us start with some preliminary observations, which will allow us to assume some extra properties of the ground state without loss of generality. Fix Λ and let $\{M_i\}_i$ a decomposition of the identity on $\bar{\Lambda}$ into orthogonal projectors, such that $[M_i, \mathbf{H}_{\bar{\Lambda}}] = [M_i, \Pi_{\text{TC},\Lambda}^\perp] = 0$ for all i . We want to show that is sufficient to study ω restricted to the subspace corresponding to each M_i . This is the content of the following lemma.

Lemma 4. *Let ω be a ground state. Fix Λ and let $\{M_i\}_i$ as above. Whenever $\omega(M_i) \neq 0$, denote by*

$$\omega_i(A) = \frac{\omega(M_i A M_i)}{\omega(M_i)}.$$

Then ω_i is also a ground state. Moreover, if for every i it holds that $\omega_i(\Pi_{\text{TC},\Lambda}^\perp) = 0$, then also $\omega(\Pi_{\text{TC},\Lambda}^\perp) = 0$.

Proof. ω_i is clearly a positive linear functional on local observables so that $\omega_i(\mathbb{1}) = 1$. It can then be extended to a state on \mathcal{A} . The fact that M_i commutes with the Hamiltonian makes ω_i fulfil trivially eq. (11), so it is a ground state. Finally, we observe that

$$\begin{aligned} \omega(\Pi_{\text{TC},\Lambda}^\perp) &= \text{tr}(\rho_{\bar{\Lambda}} \Pi_{\text{TC},\Lambda}^\perp) = \sum_i \text{tr}(M_i \rho_{\bar{\Lambda}} \Pi_{\text{TC},\Lambda}^\perp M_i) = \\ &= \sum_i \text{tr}(M_i \rho_{\bar{\Lambda}} M_i \Pi_{\text{TC},\Lambda}^\perp) = \sum_i \omega(M_i) \omega_i(\Pi_{\text{TC},\Lambda}^\perp), \end{aligned}$$

so that the last claim of the lemma follows. \square

We will use such lemma to make two extra assumptions. The first one allows to assume that the ground state is supported, in each site, only in one of the two subspaces (TC or tiling). For that, given a finite region $R \subset \mathbb{Z}^2$ we consider signatures $\sigma = (\sigma_i)_{i \in R}$ where each $\sigma_i \in \{\text{TC}, \text{tiling}\}$. We denote by P_σ the projector onto the set of states of signature σ . It is easy to see that they satisfy the condition of lemma 4. The second assumption is that $\rho_{\bar{\Lambda}}$ commutes with the Toric Code stabilizers. Again, it is sufficient to consider the projectors onto the eigenspaces of such stabilizers, and the result follows from lemma 4.

As a second step, we will show that for any ground state for which a square boundary is completely supported in the TC subspace, the interior will be as well; for this we will assume that all square regions have smooth edges as in fig. 5.

Lemma 5. *Take two concentric square regions $\Lambda' \subsetneq \Lambda$, and a ground state ω of $\mathbf{H}^{(d)}$ with a signature σ on $\bar{\Lambda}$. Assume that $\sigma_s = \text{TC}$ on all sites s of $\partial\Lambda' \subset \Lambda \setminus \Lambda'$. Moreover, assume that $\rho_{\bar{\Lambda}}$ commutes with the Toric Code stabilizers that couple Λ' with $\partial\Lambda'$. Then $\sigma_s = \text{TC}$ all sites $s \in \Lambda'$.*

Proof. Denote with $T \subseteq \Lambda'$ the set of all sites $\sigma \in \Lambda'$ that satisfy $\sigma = \text{tiling}$.

Consider the CPTP map Φ_1 acting on T that on all those sites, traces out the tiling sector and replaces it with the maximally mixed state on the TC subspace, i.e.

$$\Phi_1(\rho) = \text{tr}_T(\rho) \otimes \left(\frac{\mathbb{1}_T^{(\text{TC})}}{\text{tr} \mathbb{1}_T^{(\text{TC})}} \oplus 0_T^{(\text{tiling})} \right).$$

Let us now consider a map Φ_2 , acting on $\bar{\Lambda}'$, which implements the following operations: first measures the Toric Code projectors which overlap with Λ' , and then, conditioned on the syndrome of the measurement, applies a unitary operator which corrects as many as code errors as possible [30]. This can be constructed by choosing as Kraus operators of Φ_2 the product of the projector onto the different syndrome subspaces multiplied on the left with the corresponding unitary operator. We extend this map on the tiling subspace with the identity map, in order to make it a CPTP map. Then eq. (13) implies that

$$\text{tr}(\rho_{\bar{\Lambda}} \mathbf{H}_{\bar{\Lambda}}) \leq \text{tr}(\Phi_2 \circ \Phi_1(\rho_{\bar{\Lambda}}) \mathbf{H}_{\bar{\Lambda}}). \quad (14)$$

We now consider $\tilde{\rho}_{\bar{\Lambda}} = \Phi_2 \circ \Phi_1(\rho_{\bar{\Lambda}})$. For any \mathbf{h} whose support is disjoint from Λ' , since $\tilde{\rho}_{\bar{\Lambda}'}$ has the same reduced density matrix as $\rho_{\bar{\Lambda}'}$ outside of Λ' , we have that $\text{tr}(\rho_{\bar{\Lambda}} \mathbf{h}) = \text{tr}(\tilde{\rho}_{\bar{\Lambda}} \mathbf{h})$. Thus eq. (14) reduces to $\text{tr}(\rho_{\bar{\Lambda}} \mathbf{H}_{\bar{\Lambda}'}) \leq \text{tr}(\tilde{\rho}_{\bar{\Lambda}} \mathbf{H}_{\bar{\Lambda}'})$, where in $\mathbf{H}_{\bar{\Lambda}'}$ only appears Hamiltonian terms whose support intersects Λ' . To finish the proof, we need to find a contradiction assuming that T is not empty. First of all, notice that $\tilde{\rho}_{\bar{\Lambda}}$ is completely supported on the TC subspace in $\bar{\Lambda}'$, and that can violate at most 2 of the Toric Code Stabilizers (at most one plaquette and one star operator, since any pair of violation would have been destroyed by the action of Φ_2). So $\text{tr}(\tilde{\rho}_{\bar{\Lambda}} \mathbf{H}_{\bar{\Lambda}'})$ can at most be equal to 2. On the other hand, even with the bonus gained in bottom-left corners of regions supported on the tiling subspace (i.e. the bonus of 1/2 for the all-black tile used to resolve the ground state degeneracy for the periodic tiling pattern), the penalties coming from mixed signatures in $\text{tr}(\rho_{\bar{\Lambda}} \mathbf{H}_{\bar{\Lambda}'})$ are higher (an overall penalty of at least 7/2 for each mismatch). \square

In the next lemma, we generalise the previous one for the case in which some sites on $\partial\Lambda'$ are in the tiling sector.

Lemma 6. *Take two concentric square regions $\Lambda' \subsetneq \Lambda$, and a ground state ω of $\mathbf{H}^{(d)}$ with a signature σ on $\bar{\Lambda}$.*

$\bar{\Lambda}$. Moreover, assume that $\rho_{\bar{\Lambda}}$ commutes with the Toric Code stabilizers that couple Λ' with $\partial\Lambda'$. Let α be the number of sites $s \in \partial\Lambda'$ for which $\sigma_s = \text{tiling}$, and β the sum of signature mismatches within Λ' —i.e. the number of neighbouring $s, s' \in \Lambda$ for which $\sigma_s \neq \sigma_{s'}$ —plus the number of period markers within Λ' . Then $\beta \leq \frac{4}{7}(1 + 4\alpha)$.


Proof. We follow the same procedure as in the proof for lemma 5, obtaining a new state $\tilde{\rho}_{\bar{\Lambda}}$ on $\bar{\Lambda}$, such that

$$\text{tr}(\rho_{\bar{\Lambda}} \mathbf{H}_{\bar{\Lambda}}) \leq \text{tr}(\tilde{\rho}_{\bar{\Lambda}} \mathbf{H}_{\bar{\Lambda}}). \quad (15)$$

Again, let $T \subset \Lambda'$ the set of sites with tiling signature. Let us consider now the interactions \mathbf{h} in $\mathbf{H}_{\bar{\Lambda}}$ and compare the values $\text{tr}(\tilde{\rho}_{\bar{\Lambda}} \mathbf{h})$ and $\text{tr}(\rho_{\bar{\Lambda}} \mathbf{h})$. As in previous lemma, if \mathbf{h} do not overlap with Λ' , then $\text{tr}(\tilde{\rho}_{\bar{\Lambda}} \mathbf{h}) = \text{tr}(\rho_{\bar{\Lambda}} \mathbf{h})$. Since $\tilde{\rho}_{\bar{\Lambda}}$ is in the TC subspace, and can violate at most 2 stabilizers, its energy can be at most $2 + 8\alpha$ (the signature in $\partial\Lambda'$ has not changed, and each spin in the tiling subspace can violate up to 4 Ising-type interactions). On the other hand, since there are at least β signature mismatches for $\rho_{\bar{\Lambda}}$, and each of them has an energy of at least $7/2$ (again, this is lower than 4 because of the $1/2$ bonus given to the all-black tile), we have that $\text{tr}(\rho_{\bar{\Lambda}} \mathbf{H}_{\bar{\Lambda}}) \geq \frac{7}{2}\beta$. Inserting these two bounds into eq. (15), we obtain the desired bound. \square

In the following, we will show that if we pick the outer square in lemma 5 large enough, we are bound to find an inner concentric square—of at least a third of the outer square's size—for which we can then apply lemma 5 or lemma 6. In the pictures of the following lemma, we have coloured with black the spins which are in the TC subspace, and in yellow the ones which are not.

Lemma 7. Take some square Λ of side length $3k$, where $k = 10^6 N_d^2$ and consider a ground state ω of $\mathbf{H}^{(d)}$ with a signature σ on $\bar{\Lambda}$. Subdivide Λ into $k \times k$ squares:



Then $\sigma_i = \text{TC}$ for all i in the centre C .


Proof. We start with a few preliminary observations, and we refer the reader to fig. 5 for verification. The boundary $\partial\Lambda$ contains precisely $3 \times 3k$ spins on each side, and thus $12 \times 3k = 36k$ spins overall. We denote this spin count with $|\partial\Lambda|$ in the context of this proof. By lemma 6 (assuming that for all $s \in \partial\Lambda : \sigma_s = \text{tiling}$), we thus know that we can have at most $\frac{4}{7}(1 + 4 \times |\partial\Lambda|) < 83k$ penalties from signature mismatches or period markers within Λ .

The overall area of Λ encompasses $9k^2$ tiles, and we count $|\Lambda| = 2 \times 9k^2 + 2 \times 3k$ spins. Every sub-square of size $N_d \times N_d$ which is not fully in the TC subspace carries a penalty ≥ 1 —note that this holds true regardless of the bonus terms present in the tiling, as the period penalty and TC-tiling mismatch are both larger. This means that at most a fraction of

$$\frac{83k}{|\Lambda|/N_d^2} = \frac{83N_d^2}{18k + 6} = \frac{83}{18 \times 10^6 + 6/N_d^2} < \frac{1}{10000}$$

of spins $s \in \Lambda$ can have signature $\sigma_s = \text{tiling}$. For a subset $A \subset \Lambda$, we denote this fraction with $f(A)$.

So let us assume that $f(\Lambda) < 1/10000$. Take Λ and shrink it uniformly by at most $k/10$, by which we mean we shrink the square on each side by one tile at a time, i.e. while keeping smooth boundaries as in fig. 1. This sweeps a region A which covers at least $1/10$ of the area of Λ , and thus in particular the number of spins $|A| \geq |\Lambda|/10$. Since $f(\Lambda) < 1/10000$, it follows that $f(A) < 1000$. This immediately implies that there exists a square $\Lambda' \subset A$ —concentric with Λ —which satisfies $f(\partial\Lambda') < 1/1000$.



Two things may happen. If in between the centre square C and Λ' there exists another square Λ'' (i.e. with $\partial\Lambda'' \subset \Lambda' \setminus C$) such that for any $s \in \partial\Lambda''$, we have $\sigma_s = \text{TC}$, then lemma 5 immediately implies that $\sigma_i = \text{TC}$ for all i in the centre C , and the claim follows.

It remains to analyse the case where no such square Λ'' exists. We first apply lemma 6 again, this time to Λ' : as $f(\partial\Lambda') < 1/1000$, and $|\partial\Lambda'| \leq |\partial\Lambda| \leq 36k$, we know that there exist at most $36k/1000 \times 8 \leq k/3$ spins $s \in \Lambda'$ with $\sigma_s = \text{tiling}$. But since no Λ'' exists with a boundary $\partial\Lambda''$ completely with tiling signature, there have to be at least $k - k/10 = 9k/10$ spins within Λ' with a tiling signature (one within the boundary for each concentric square between Λ' and C). Contradiction, and the claim follows. \square

All the above results lead trivially to

Corollary 8. All ground states of $\mathbf{H}^{(d)}$ are fully supported in the TC subspace.