

## Solving frustrated Ising models using tensor networks

Bram Vanhecke,<sup>1,\*</sup> Jeanne Colbois<sup>2,†</sup>, Laurens Vanderstraeten,<sup>1</sup> Frank Verstraete,<sup>1</sup> and Frédéric Mila<sup>2</sup><sup>1</sup>*Department of Physics and Astronomy, University of Ghent, Krijgslaan 281, 9000 Gent, Belgium*<sup>2</sup>*Institute of Physics, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*

(Received 7 July 2020; revised 9 December 2020; accepted 14 December 2020; published 13 January 2021)

Motivated by the recent success of tensor networks to calculate the residual entropy of spin ice and kagome Ising models, we develop a general framework to study frustrated Ising models in terms of infinite tensor networks that can be contracted using standard algorithms for infinite systems. This is achieved by reformulating the problem as local rules for configurations on overlapping clusters chosen in such a way that they relieve the frustration, i.e., that the energy can be minimized independently on each cluster. We show that optimizing the choice of clusters, including the weight on shared bonds, is crucial for the contractibility of the tensor networks, and we derive some basic rules and a linear program to implement them. We illustrate the power of the method by computing the residual entropy of a frustrated Ising spin system on a kagome lattice with next-next-nearest-neighbor interactions, vastly outperforming Monte Carlo methods in speed and accuracy. The extension to finite temperatures is briefly discussed.

DOI: [10.1103/PhysRevResearch.3.013041](https://doi.org/10.1103/PhysRevResearch.3.013041)

## I. INTRODUCTION

One of the most beautiful manifestations of emergent behavior in statistical physics can be found in the arena of frustrated spin systems [1]. Frustration in a classical spin system occurs whenever it is impossible to find a spin configuration which minimizes each and every term of the Hamiltonian simultaneously, leading to macroscopic ground-state degeneracies and giving rise to interesting zero-temperature physics such as effective realizations of gauge theories [2].

Early exact results in this context were obtained for antiferromagnetic nearest-neighbor Ising models on triangular and kagome lattices [3,4] using Kauffman and Onsager's method, for frustrated Ising models on all planar two-dimensional lattices with nearest-neighbor interactions using a mapping to free fermions [5,6], and for more general systems such as planar spin ice [7] using Bethe ansatz techniques [8].

It has, however, proven difficult to treat frustration in generic (i.e., nonintegrable) models: to reach the low-energy phase space and sample it efficiently, Monte Carlo methods require *ad hoc* nonlocal cluster updates to fight both critical slowing-down [9,10] and frustration [11,12]. In addition, calculating the free energy requires the use of thermodynamic integration, making zero-temperature residual entropies hard to determine accurately.

Tensor networks [13] provide a new computational approach for studying ground states of classical lattice models with strong correlations, as recently demonstrated by

the determination of the residual entropy of ice and dimer models in three-dimensional lattices with unprecedented precision [14]. There is some freedom when expressing a partition function as a tensor network. The formulation of the tensor network in Ref. [14] relies on preexisting knowledge of ground-state *local rules*—rules that all the ground-state configurations of a model must satisfy—easily implemented at the level of the tensor.

In this paper, we generalize the applicability of tensor networks to more generic frustrated spin systems. For this, we first revisit nearest-neighbor frustrated systems. We argue that, for a given partition function, the choice of the tensor network expression affects the convergence of the contraction algorithms, and we show that taking the zero-temperature limit of the standard formulation [15,16] is not always an option. We observe that a formulation relying on ground-state local rules, turning the computation of the partition function into a tiling problem, appears to be crucial for the contraction to converge. Furthermore, the tiles need to be selected with care, and we illustrate how this can be done. We introduce a linear program (expanding on Ref. [17]) to systematically look for such local rules in generic further-neighbor models, and we show how this yields a natural expression for the tensor network enabling the study of the full ground-state manifold. The construction holds the roots of a generalization to finite temperature. To demonstrate the power of the method, we apply it to a frustrated Ising model with further-neighbour couplings on a kagome lattice, obtain the residual entropy with a very high accuracy, and use the local rules to make some exact statements regarding the physics of the ground-state manifold.

## II. STANDARD CONSTRUCTION

Partition functions for classical spin systems can be expressed as contractions of tensor networks in the spirit of the transfer matrix formalism, a representation which is not unique. The standard construction consists in associating with

\*bavhecke.vanhecke@ugent.be

†jeanne.colbois@epfl.ch

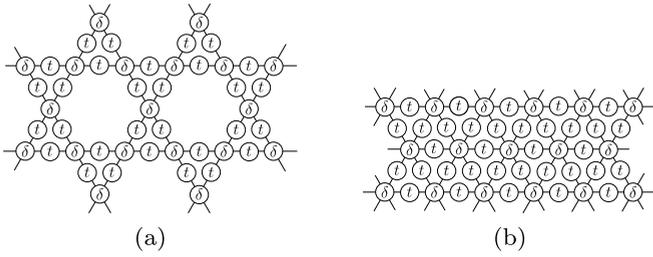


FIG. 1. Standard tensor network construction for the partition function on triangular and kagome lattices. The matrices  $t$  carry the Boltzmann weights, while the  $\delta$  tensors enforce that neighboring  $t$ 's share the same spin.

each interaction a matrix  $t$  accounting for its Boltzmann weight and with each on-site variable a  $\delta$  tensor, i.e., a tensor whose rank corresponds to the number of interactions involving that site and which is 1 only when all of its indices take the same value. In the particular case of kagome and triangular lattice Ising (anti)ferromagnets, this standard formulation leads to the tensor networks in Fig. 1, where the Boltzmann weights are given by

$$t^{\sigma_i, \sigma_j} = e^{-\beta J \sigma_i \sigma_j}, \quad (1)$$

so that matrix  $t$  reads

$$t = \begin{pmatrix} e^{-\beta J} & e^{+\beta J} \\ e^{+\beta J} & e^{-\beta J} \end{pmatrix}, \quad (2)$$

with  $J > 0$  for the antiferromagnet. Contracting the tensor network amounts to finding the leading eigenvalue and leading eigenvector of the row-to-row transfer matrices (see, e.g., [16]). When the algorithm converges, the logarithm of the leading eigenvalue, directly related to the free energy per site at the given inverse temperature  $\beta$ , is obtained with an extremely high accuracy.

#### Issue with the zero-temperature limit

However, it is obvious that low-temperature properties cannot be directly probed from the standard construction since the zero-temperature limit of Eq. (2) cannot be taken. Similarly, the partition function is ill defined in that limit. In simple cases, this problem can be solved by considering the regularized partition function  $\mathcal{Z}_0$  whose zero-temperature limit is always well defined and is directly related to the residual entropy,

$$\mathcal{Z}_0 := e^{\beta E_0 N} \mathcal{Z}, \quad S = \lim_{\beta \rightarrow \infty} \frac{1}{N} \ln(\mathcal{Z}_0), \quad (3)$$

where we have introduced the ground-state energy per site  $E_0$  and the number of sites  $N$ .

Indeed, to compute  $\mathcal{Z}_0$  instead of  $\mathcal{Z}$ , one has to contract the tensor network based on  $t_0^{\sigma_i, \sigma_j} = e^{\beta \frac{E_0}{z}} t^{\sigma_i, \sigma_j}$ , where  $z$  is the number of bonds per site. In a nonfrustrated system, all the pair interactions are minimized simultaneously, removing all exponentially diverging matrix elements, hence ensuring that the zero-temperature limit can be taken in both  $t_0$  and  $\mathcal{Z}_0$ . However, in a frustrated system, the pair interactions cannot be minimized simultaneously and  $t_0$  still contains exponentially diverging factors. For instance, compare the tensors for

the ferromagnetic and antiferromagnetic Ising models on a kagome lattice [18]:

$$t_0^F = \begin{pmatrix} 1 & e^{-2\beta|J|} \\ e^{-2\beta|J|} & 1 \end{pmatrix}, \quad t_0^{AF} = \begin{pmatrix} e^{-\frac{4}{3}\beta J} & e^{\frac{2}{3}\beta J} \\ e^{\frac{2}{3}\beta J} & e^{-\frac{4}{3}\beta J} \end{pmatrix}. \quad (4)$$

### III. NEAREST-NEIGHBOR ISING ANTIFERROMAGNET ON A KAGOME LATTICE

#### Ground state

To build up to a generic construction, we start by reexploring the simple frustrated models where a solution is known. The simplest is the kagome lattice Ising antiferromagnet. In this model, the ground-state configurations have to satisfy a “two-up one-down, two-down one-up” rule (no ferromagnetic triangles). This is easily seen by writing the Hamiltonian as a sum of triangular Hamiltonians:

$$H = \sum_{\langle i,j \rangle} J \sigma_i \sigma_j = \sum_{\substack{\Delta_{i,j,k} \\ \nabla_{i,j,k}}} J (\sigma_i \sigma_j + \sigma_j \sigma_k + \sigma_k \sigma_i) \quad (5)$$

$$:= \sum_{\Delta_{i,j,k}} H_{\Delta_{i,j,k}} + \sum_{\nabla_{i,j,k}} H_{\nabla_{i,j,k}}. \quad (6)$$

One triangular Hamiltonian is minimized by nonferromagnetic spin configurations on the triangle. Since the triangular Hamiltonians can be simultaneously minimized, *all* the ground states of the model can be described as tilings of two-up one-down, two-down one-up triangles on a kagome lattice, where the triangular tiles fit if the shared spin is the same.

This is easily translated into a tensor network (slightly different from the one in Ref. [14]) on a (dual) honeycomb lattice. The prescription is as follows (Fig. 2):

(1) At each site of the dual lattice (center of the kagome triangles), place a  $\delta$  tensor of rank 3 and bond dimension 6 describing the six ground-state configurations of this triangle.

(2) On each bond of the dual lattice (sites of the kagome lattice), place a bond matrix  $P$  with bond dimension 6, which is 1 if the two connected tensors assign the same value to their shared spin and 0 otherwise.

(3) Reduce the bond dimension of the tensor network to 2 by performing a singular value decomposition (SVD) on the  $P$  tensors and grouping the resulting tensors with the  $\delta$  tensors on the triangles.

This tensor network is well defined and provides the ground-state entropy of the kagome lattice with a precision of  $10^{-10}$  (Table I). This example demonstrates that it can be very useful to use clusters (here triangles) to build a tensor network. However, the choice of clusters, which is rather natural in the case of the kagome lattice, is in general a subtle issue, as we now show with the example of the triangular lattice.

### IV. NEAREST-NEIGHBOR ISING ANTIFERROMAGNET ON A TRIANGULAR LATTICE

#### A. Ground state

We proceed with the archetype of frustration: the triangular lattice Ising antiferromagnet. Inspired by the kagome construction, we look for *tiles* that can be used to build *all*

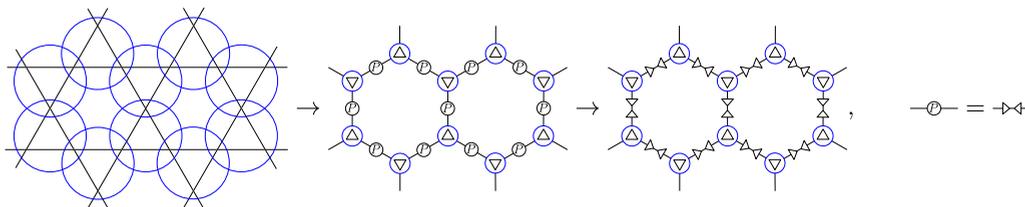


FIG. 2. Tensor network construction for the ground state of the kagome lattice Ising antiferromagnet: the Hamiltonian can be split up into triangular terms. The configurations which minimize the triangular Hamiltonian can be tiled to create ground states. We can build a tensor network to count the tilings and study the ground-state manifold by associating a tensor with each triangle and introducing bond matrices  $P$  enforcing that spins must match. The bond dimension is significantly reduced by performing an SVD on  $P$ . At finite temperature, the construction is the same, only the tensors are promoted to consider more configurations and provide them with a Boltzmann weight.

the ground states. We want to find them as local ground-state configurations of *local* Hamiltonians (the equivalent of the triangular Hamiltonian in the previous section) which can be *simultaneously minimized*. The latter criterion is essential to ensure that *each* ground state can be described using these tiles.

For this, we note that the Hamiltonian can be written as a sum of terms acting only on one type of triangle, for instance,  $\Delta$  triangles:

$$H = \sum_{\langle i,j \rangle} J\sigma_i\sigma_j = \sum_{\Delta_{i,j,k}} J(\sigma_i\sigma_j + \sigma_j\sigma_k + \sigma_k\sigma_i). \quad (7)$$

The triangular Hamiltonian is minimized by nonferromagnetic triangles. Since there exists a state which minimizes all the triangular Hamiltonians, the set of all ground states can be obtained by tiling nonferromagnetic down triangles, which fit if the spins in the overlap of three triangles have the same orientation [19]. The corresponding tensor network has  $\delta$  tensors at the centers of up triangles and rank 3  $P$  tensors enforcing the consistency of the spin shared by three  $\delta$  tensors (Fig. 3).

Another valid splitting of the Hamiltonian is to share bonds between up and down triangles:

$$H = \sum_{\langle i,j \rangle} J\sigma_i\sigma_j = \sum_{\substack{\Delta_{i,j,k} \\ \nabla_{i,j,k}}} \frac{J}{2} (\sigma_i\sigma_j + \sigma_j\sigma_k + \sigma_k\sigma_i). \quad (8)$$

This splitting amounts to tiling nonferromagnetic up *and* down triangles with the condition that, on a shared bond, the two spins must match; the corresponding tensor network is defined on the honeycomb lattice and has  $\delta$  tensors on up *and* on down triangles, with bond matrices  $P$  now taking care of two spins (Fig. 3). Note that we have chosen to give the tensor network minimal connectivity: not all triangles that share a

spin are connected; some shared spins are implicitly enforced to be the same via multiple bonds.

These two splittings are equally valid (in the thermodynamic limit or with periodic boundary conditions), and both solve the regularization problem by working directly in the ground state. However, standard contraction algorithms fail to converge for the first construction, while they converge and lead to the correct answer for the second one (Fig. 4, Table I).

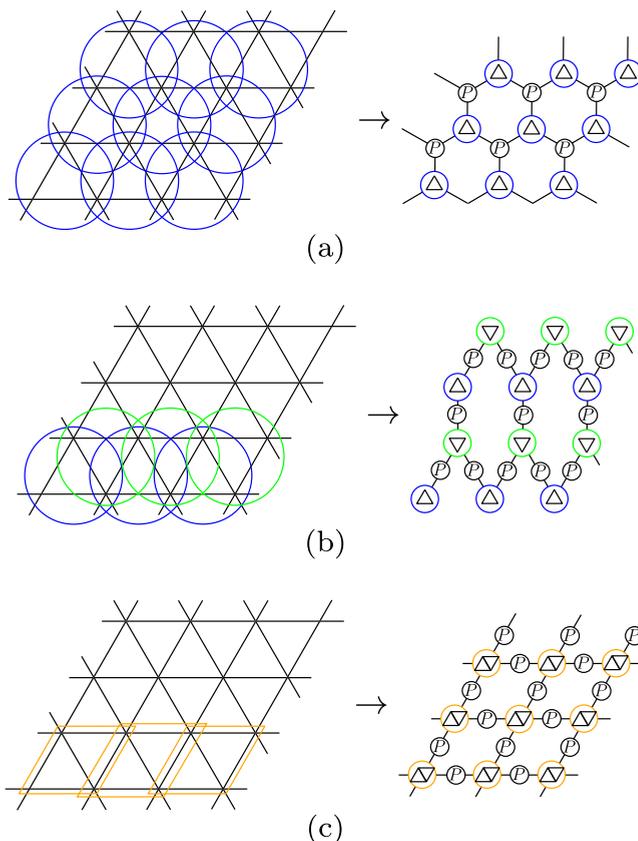


FIG. 3. Tensor network construction for the ground state of the triangular lattice Ising antiferromagnet: the Hamiltonian can be split up into triangular terms in various ways. The configurations which minimize the triangular Hamiltonian can be tiled to create ground states. Top: Construction corresponding to Eq. (7). Middle: Construction corresponding to Eq. (8). Bottom: Construction based on the Hamiltonian tessellation, Eq. (9). The bond dimension is significantly reduced by performing an SVD on the bond matrices  $P$ .

TABLE I. Tensor-network results obtained using vumps on the row-to-row transfer matrix.  $D$  is the MPS bond dimension. Taken from Ref. [14].

	AF Ising	
	On kagome lattice	On triangular lattice
MPS	0.5018331646 ( $D = 10$ )	0.3230659407 ( $D = 250$ )
Exact	0.5018331646	0.3230659669

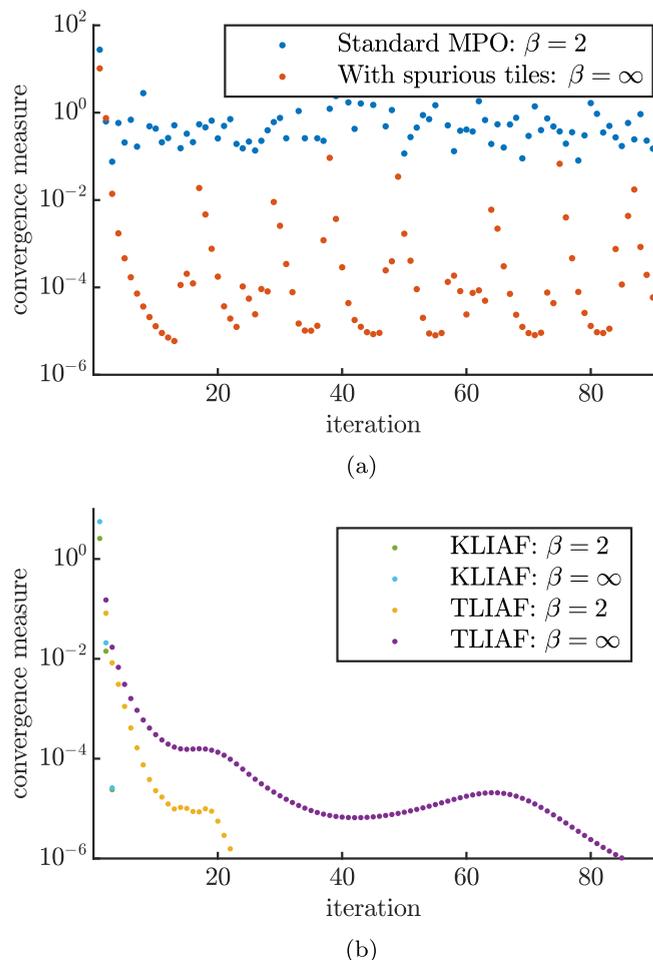


FIG. 4. (a) Convergence of the vumps algorithm [21–23] for the antiferromagnetic triangular lattice Ising model at  $\beta = 2$  using the standard tensor network from Fig. 1 and at  $\beta = \infty$  using the construction based on Eq. (7). (b) Vumps convergence for the antiferromagnetic kagome lattice Ising model at  $\beta = 2$  and  $\beta = \infty$  and the same for the antiferromagnetic triangular Ising model using the construction of Eq. (9). We use MPS with a bond dimension of  $\chi = 80$  and a variational convergence measure; see Ref. [23]. We observed similar behavior using the corner transfer matrix renormalization-group algorithm [15,24], and a similar issue was observed for real-space renormalization techniques in Ref. [25].

The main difference between the two cases is that, while in the second construction the constraint forbidding down triangles to be ferromagnetic is imposed at the level of the tensors, in the first construction it is imposed nonlocally. Indeed, since the Hamiltonians in Eqs. (7) and (8) are the same, they must have the same energy for all global states, implying that if a state contains a down triangle which is ferromagnetic, it must also contain a ferromagnetic up triangle. The key point is that these two triangles can be arbitrarily far apart. Accordingly, approximate algorithms, which are based on large but finite bond dimensions, fail to converge.

### B. Hamiltonian tessellation

In general, we do not have such insight into the ground state of frustrated models. To generalize tensor network

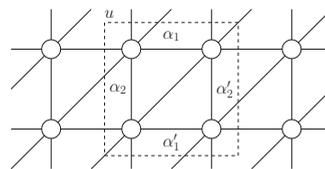
constructions, the question thus boils down to being able to select *a priori* among all possible splittings of the Hamiltonian of interest, the equivalent of Eq. (8) and not Eq. (7). Let us see how this can be done for the triangular lattice.

These two splittings can be seen as instances of the following generic *Hamiltonian tessellation* (we use this term to describe a set of ways of splitting the Hamiltonian):

$$\mathcal{H} = \sum_{c \in \mathcal{T}_u} \sum_{n \in c} \alpha_n^c h_n = \sum_{c \in \mathcal{T}_u} H_c^{(\alpha)}, \quad (9)$$

$$\sum_{c \in \mathcal{T}_u | n \in c} \alpha_n^c = 1, \quad \forall n, \quad (10)$$

where for later convenience we have considered a cluster  $u$  regrouping two triangles:



$\mathcal{T}_u$  is the set of clusters obtained from translating  $u$  on the lattice (with overlaps). Each cluster is seen as a collection of bonds (indexed by  $n$ ), and with each bond Hamiltonian  $h_n$  ( $h_{(i,j)} = J\sigma_i\sigma_j$ ) we associate weights  $\alpha_n^c$  specifying how much of it is accounted for in each cluster [Eq. (10) imposes that terms appearing in a single cluster have weight 1]. In the following, to ensure translation invariance, we always choose the  $\alpha_n^c$  to be the same for each  $c \in \mathcal{T}_u$  [20]. In our triangular case, each cluster has five interaction terms, four of which are shared with neighboring clusters, and the associated weights must satisfy the constraints

$$\alpha'_{1,2} = 1 - \alpha_{1,2} \quad (11)$$

by translation invariance and Eq. (10).

Remember that the aim is to find tiles to build all the ground states. We find these tiles as local ground-state configurations  $\{C_u^{(\alpha)}\}$  on  $u$  minimizing the local Hamiltonian  $H_u^{(\alpha)}$ . Depending on the weights, we get different ground-state configurations  $\{C_u^{(\alpha)}\}$ ; only the weights for which all the local Hamiltonians can be simultaneously minimized provide tiles which can be used to describe the whole ground-state manifold. In our case, this further restricts the weights to

$$\alpha_1 = \alpha_2 \in [0, 1]. \quad (12)$$

This is the (convex) set of weights which satisfy

$$H^{(\alpha)}(C) \geq -J \forall C \quad \text{on } u, \quad (13)$$

where  $-J$  is the ground-state energy per cluster. The boundaries of the convex set are defined by some of these inequalities becoming equalities.

In this formulation, we can see in a new light what happens on the triangular lattice. By construction, for any weights in the convex set defined by Eqs. (11) and (12), all the ground states can be constructed as tilings of the local ground-state configurations. On the one hand, Eq. (8) corresponds to Eq. (9) with  $\alpha_1 = \alpha_2 = 1/2$ . There are 10 local ground-state configurations of the unit  $u$ , hence 10 tiles. These are the

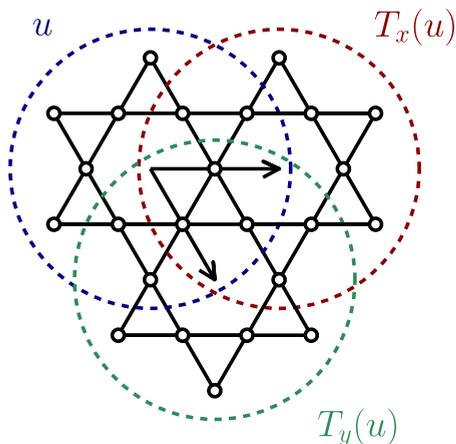


FIG. 5. A tessellation of spins on a kagome lattice: the cluster  $u$  consists of 12 spins, and it shares 5 spins with each of the translated clusters  $T_x(u)$  and  $T_y(u)$ .

configurations containing no ferromagnetic triangles. On the other hand, Eq. (7) corresponds to Eq. (9) with  $\alpha_1 = \alpha_2 = 0$ . At this point (which lies on the boundary of the convex set), an additional accidental ground-state degeneracy occurs: configurations for which the up triangle is ferromagnetic now have the ground-state energy as well. These two additional tiles cannot play a role in the ground-state manifold, since there are weights for which they are not ground-state tiles; so, they cannot fit into any global ground state. We call such tiles *spurious* because they do not really belong to the ensemble of ground-state tiles. Thus, according to our observation that contraction is possible for the tessellation of Eq. (8) but not for that of Eq. (7), it sounds like a good idea to get rid of such tiles to ensure the convergence of the tensor network.

## V. GENERIC IMPLEMENTATION

### A. Maximal lower bound

All of the above can be quite straightforwardly adapted for a generic system of  $d$ -level spins  $s_i$  on a lattice, with a translation invariant Hamiltonian  $\mathcal{H}$  containing only local interaction terms  $h_n$  of strictly bounded range (e.g., finite-range further-neighbor pair interactions),

$$H(\{\sigma\}) = \sum_n h_n(\{\sigma\}_n), \quad (14)$$

where  $\{\sigma\}_n$  denotes the subset of spins taking part in interaction  $n$ . Given a reference cluster of spins  $u$ , we cover the lattice with the set  $\mathcal{T}_u$  of overlapping translated  $u$ 's such that the Hamiltonian can be rewritten as a sum of strictly local terms acting within a single cluster (see, for instance, Fig. 5). Just as in the triangular case, we associate with each Hamiltonian term  $h_n$  weights  $\alpha_n^c$  describing how they are shared between clusters, recovering exactly the expression that we gave for the triangular lattice in Eqs. (9) and (10). Since, in this form, the Hamiltonian contains only terms that act within a cluster, the minimum of  $H_u^{(\alpha)}$  with respect to the spin configurations of  $u$  implies a lower bound on the global ground-state energy. This bound can be optimized by maximizing over  $\alpha_n^u$

[17,26–28]. This optimization (which was, in particular, leveraged in Ref. [17] to obtain ground-state energies of generalized Ising models) can be expressed as a linear program [29],

$$E_u \leftarrow \max_{\alpha} E, \quad \text{with} \quad \begin{cases} H_u^{(\alpha)}(C) \geq E \quad \forall C, \\ \sum_{c \in \mathcal{T}_u | n \in c} \alpha_n^c = 1, \end{cases} \quad (15)$$

where the result of the maximization,  $E_u$ , is the candidate ground-state energy per cluster and where  $C$  goes through the configurations of  $u$ . This program may be solved using a standard linear programming toolbox.

### B. Getting the most out of knowing the ground-state energy

Our point is that if the maximal lower bound is saturated (i.e., if  $u$  is such that there exists a global state which has  $E_u$  as the energy per cluster or, equivalently, when this lower bound matches an upper bound), one gets more than just the ground-state energy. Indeed, for the weights  $\{\alpha_n^u\}$  which realize this maximal lower bound, by construction the local cluster Hamiltonians  $H_u^{(\alpha)}$  can be simultaneously minimized if and only if the maximal lower bound is saturated. We say that such a Hamiltonian has *minimal frustration*. In this case, all the ground states are characterized as tilings of the configurations of the cluster  $u$  belonging to the set

$$G^{(\alpha)} := \{C | H_u^{(\alpha)}(C) = E_u\}. \quad (16)$$

In this sense, the set of tiles  $G^{(\alpha)}$  is a local rule.

There are, however, many solutions of (15), namely, all sets of weights  $\{\alpha_n^u\}$  satisfying

$$H_u^{(\alpha)}(C) \geq E_u \quad \text{for all configurations } C \text{ of } u. \quad (17)$$

As we have seen in the triangular lattice case, not all  $\{\alpha_n^u\}$  will do. In the space of the weights, the set of  $\{\alpha_n^u\}$  for which all these inequalities are satisfied takes the form of a convex set, which we refer to as  $A^u$ , corresponding to the generalization of Eq. (12). The set of ground-state tiles  $G_u^{(\alpha)}$  does not depend on the weights in the *interior* of  $A^u$ . However, just like in the triangular lattice case, the boundary is defined by some of the inequalities becoming equalities, and accidental degeneracies will occur:

$$G_u^{(\alpha) \in \text{Int}(A^u)} \subset G_u^{(\alpha) \in \text{Bound}(A^u)}. \quad (18)$$

The associated additional configurations must be spurious tiles, which could spoil the contractibility of the tensor network as well as hindering the understanding of the ground-state manifold.

Thus, for a generic problem, we need

- (1) to find a cluster  $u$  such that the maximal lower bound for the ground-state energy,  $E_u$ , is saturated and
- (2) to find weights  $\{\alpha_n^u\}$  in the interior of  $A^u$ .

Note that the second step allows one to get rid of avoidable spurious tiles, but there could as well be some tiles in  $G_u^{(\alpha) \in \text{Int}(A^u)}$  which do not belong to any ground state. Getting rid of the avoidable tiles might help, but in general, because of the lack of insight into the problem, several clusters might need to be tested. Additionally, to find a point in the interior of  $A^u$ , splitting the weights evenly among clusters does not always work. The fact that the problem can be phrased as a linear program is thus very helpful: first, it allows one to

rapidly test for various candidate clusters  $u$ ; second, with a bit of additional work as mentioned below, it allows one to enforce the selection of weights in the interior of  $A''$ , thus getting rid of avoidable spurious tiles.

As a technical note, linear program solvers only output extreme points of the convex set, so simply solving (15) will systematically give  $\{\alpha''_n\}$  corresponding to avoidable spurious tiles. We show in Appendix A how to overcome this by finding boundary points that form a simplex of the same effective dimension as  $A''$ , ensuring that any point in the interior of this simplex will also lie in the interior of  $A''$  [30]. Another technical challenge is that the number of constraints scales exponentially in the number of spins per cluster. In Appendix A, we also show how to work around this problem by systematically and progressively incorporating inequalities as we build the corners of the interior simplex, such that only a very limited number of inequalities is needed.

### C. Generic tensor network

Finally, the procedure to write a contractible tensor network is easily generalized. For the ground state, the tiles  $G_u^{(\alpha) \in \text{Int}(A'')}$  are described by  $\delta$  tensors, placed on each dual vertex, coinciding with the clusters of  $\mathcal{T}_u$ . The overlapping spins matching condition is enforced by bond matrices  $P$ . Performing an SVD on the rank-deficient bond matrices keeps the tensor network bond dimension reasonably small.

### D. Testing for the saturation of the maximal lower bound

Provided that the maximal ground-state lower bound is saturated, all of the above is given. Rigorously proving that this is the case is equivalent to finding one ground state or proving that the tiles  $G_u^{(\alpha)}$  corresponding to the maximal lower bound can tile the plane. In general, this is an undecidable problem [31–34].

However, in practice, there is a whole range of models where it remains manageable, for instance, by constructing an upper bound with linear programming [17]. Moreover, the tensor network formulation typically helps to deal with this question. Indeed, if the tiles  $G_u^{(\alpha)}$  cannot tile the plane, then the associated partition function is 0 in the thermodynamic limit. Thus, reciprocally, if the (exact) leading eigenvalue associated with the transfer matrix is larger than or equal to 1, this implies that the plane can be tiled using  $G_u^{(\alpha)}$  and that the lower bound is saturated. In practice, tensor network algorithms compute the contraction of the partition function and the leading eigenvalue approximately. The convergence parameter is the bond dimension of the candidate leading eigenvector in the form of an MPS. We deduce from the above that if, for increasing MPS bond dimensions, the approximate contractions converge consistently to a leading eigenvalue which is larger than or equal to 1, we have numerical evidence that the maximal lower bound is saturated, implying in turn that we found the ground-state manifold and its degeneracy. Conversely, if the contraction does not converge, we cannot conclude: this could mean either that the lower bound is not saturated or that it is saturated but unavoidable spurious tiles spoil the convergence. In this case, one should try a different cluster.

### E. Convergence at finite temperature

Before moving to the example, we note additionally that the above construction can in principle readily be generalized to a finite-temperature tensor network. It will be the topic of another paper to show that this more generic tensor network provides accurate results and allows one to study challenging cases. Here, we just give the idea for the construction and show that it converges at finite temperatures for the triangular lattice Ising antiferromagnet.

The ground-state tensor network that we built can be seen as the zero-temperature limit of a slightly more general construction, where the  $\delta$  tensor is promoted to a tensor  $D_0$  (with a larger bond dimension) describing each configuration and its Boltzmann weight relative to the ground state. In the triangular lattice case, the finite-temperature tensor network formulation of the regularized partition function  $\mathcal{Z}_0$  associated with a Hamiltonian tessellation using  $u$  and weights  $\alpha_1 = \alpha_2 =: \alpha \in [0, 1]$  on the triangular lattice Ising antiferromagnet is thus given by bond matrices  $P$  of bond dimension 16 and by  $D_0$  tensors of the same bond dimension defined on each couple of triangles as

$$D_0^{\{\sigma\}, \{\sigma'\}, \{\sigma''\}, \{\sigma'''\}}(\alpha, \beta) = \delta^{\{\sigma\}, \{\sigma'\}, \{\sigma''\}, \{\sigma'''\}} B(\{\sigma\}, \alpha, \beta), \quad (19)$$

where the Boltzmann weight is given by

$$B(\{\sigma\}, \alpha, \beta) = e^{-\beta J[\alpha(\sigma_1\sigma_2 + \sigma_1\sigma_4) + (1-\alpha)(\sigma_2\sigma_3 + \sigma_3\sigma_4) + 2]} \quad (20)$$

and where, for short, we have denoted by  $\{\sigma\}$  the configuration of the four spins. After SVD and grouping, the bond dimension of the network is reduced to 4. Importantly, in the standard construction, tensor network algorithms fail to converge even at modest inverse temperatures, when  $\beta$  is still small enough that the values in  $t_0$  are well defined, and no “NaN” arise; the criterion for convergence is simply never met. In contrast, our tensor network can be contracted without issues at any inverse temperature (Fig. 4). The zero-temperature limit of  $D_0$  is well defined by construction, and in that limit it reduces to a  $\delta$  tensor corresponding to the 10 ground-state tiles on  $u$ .

## VI. FURTHER-NEIGHBOR ISING MODEL ON A KAGOME LATTICE

As a challenging test case, we consider a frustrated Ising model inspired by Refs. [35–39] and defined on a kagome lattice,

$$\mathcal{H} = J_1 \sum_{\langle ij \rangle} \sigma_i \sigma_j + J_2 \sum_{\langle\langle ij \rangle\rangle} \sigma_i \sigma_j + J_3 \sum_{\langle\langle\langle ij \rangle\rangle\rangle} \sigma_i \sigma_j, \quad (21)$$

where the sums run over (distance-based) first-, second-, and third-nearest neighbors, respectively, as illustrated in Fig. 6. We take  $J_1 = -1$  (ferromagnetic), and  $J_2 = J_3 = 10$  (antiferromagnetic). As our reference cluster  $u$  for the tessellation, we use a full kagome star (12 spins; Fig. 5), for which 18 weights need to be determined. From the linear program, we find a ground-state-energy lower bound  $E = \frac{2}{3}J_1 - \frac{2}{3}J_2 - J_3$  and 132 candidate ground-state tiles. The tensor network we construct for the ground-state ensemble, assuming those candidate tiles, has bond dimension 18 (very small compared to the total number of tiles in the cluster,  $2^{12} = 4096$ ). The

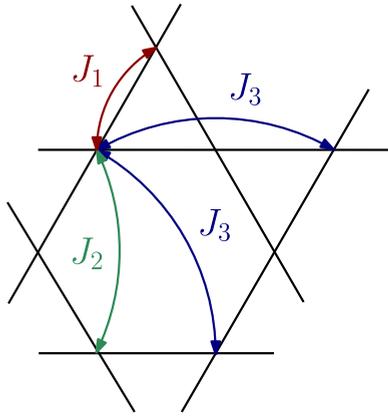


FIG. 6. The two-body Ising interactions present in test model 21.

vumps algorithm [21–23] converges nicely for this MPO for all bond dimensions of the MPS and finds a leading eigenvalue that is both real and larger than 1. We thus obtain with a good level of confidence that the ground-state tiles can tile the plane, which implies minimal frustration. The contraction—which took around 2 h on a laptop for the largest MPS bond dimension—readily provides the ground-state entropy to a very high precision (Fig. 7). Note that this method did not rely on constructing a periodic ground state or any insight from the Monte Carlo results which are presented below; the mere existence of the state in Fig. 8, however, *proves* this result.

For comparison, we calculated the residual entropy using Monte Carlo methods (the technical details are given in Appendix C). It turned out to be crucial to employ a combination of worm updates [12], single spin flip, and parallel tempering. Though one can easily generate some ground-state configurations of the model (Fig. 8), the evaluation of the residual entropy via thermodynamic integration is a huge challenge which requires thousands of CPU hours for a significantly less accurate result (compare Fig. 9 to Fig. 7).

The residual entropy obtained by contracting the tensor network is within  $10^{-7}$  of one third of the triangular lattice Ising antiferromagnet entropy, suggesting some kind of correspondence between the dominant part of the ground-state manifolds of both models. This correspondence can be under-

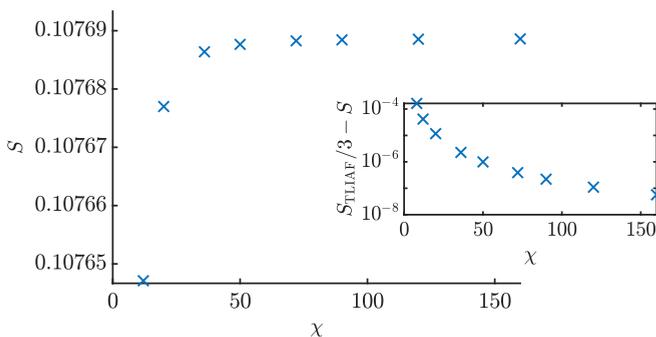


FIG. 7. The residual entropy per site of the minimally frustrated tensor network for the model in Eq. (21), obtained with the vumps algorithm with different bond dimensions  $\chi$ . Inset: The value converges to a third of the value for the Ising antiferromagnet on the triangular lattice [3].

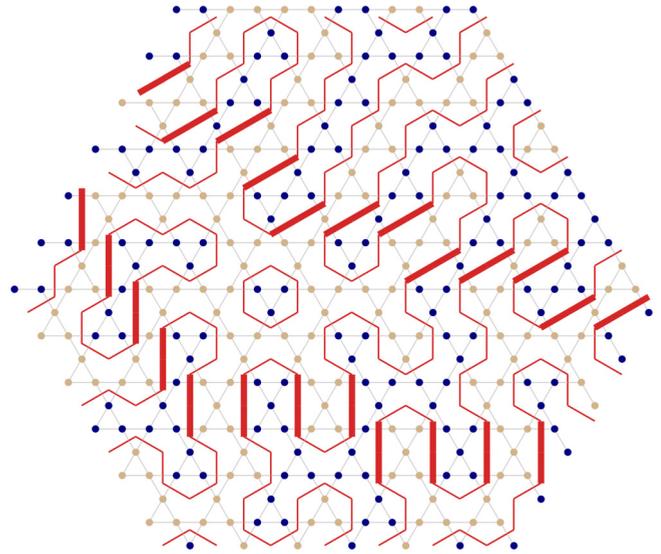


FIG. 8. An example of a spin configuration in the ground state. Red lines separate up and down spins, and red lines that cross a hexagon in a straight line are accentuated. The tiles with a thick line separate symmetry-broken sectors where all up (down) triangles are ferromagnetic. This configuration was generated during our Monte Carlo sampling and illustrates the results obtained from the tensor network.

stood thanks to exact statements based on the tiles and the tensor network construction. The 132 tiles can be split up into two types: type I tiles, for which all up (down) triangles are ferromagnetic; and type II tiles, which have one up and one down antiferromagnetic triangle (Fig. 10). To simplify the visualization, we introduce lines to separate up spins from down spins; this line must cross the hexagon with a  $120^\circ$

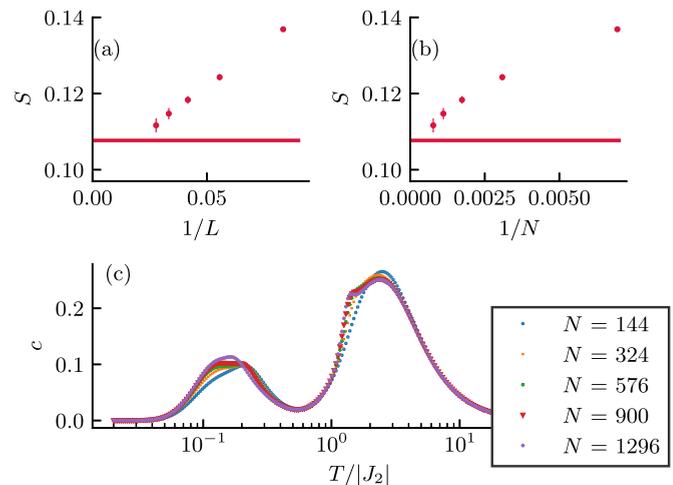


FIG. 9. Monte Carlo results for the residual entropy as a function of the inverse linear system size (a) and of the inverse system size (b). For each size, the entropy is obtained by integrating the heat capacity (c) over the temperature in the whole temperature range. The heat capacity is measured on 216 temperatures thermalized with 16 384 Monte Carlo steps (MCSs; consisting of two full updates of the state with a single spin flip, two with a dual worm, and one parallel tempering step) and measured over 1 048 576 MCSs.

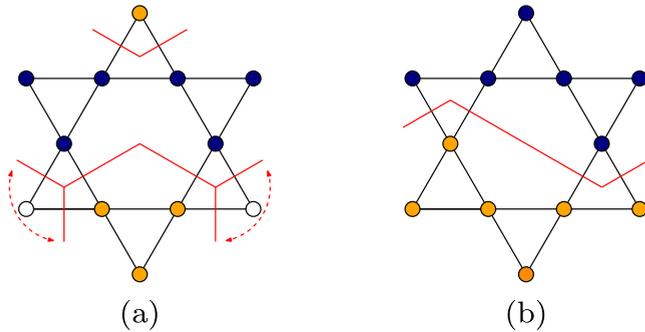


FIG. 10. Type I tiles (a) and a type II tile (b), with the line mapping. Arrows indicate that the line can lie across either of the two directions but not both. The  $J_2$  and  $J_3$  interactions are at minimal energy, and the  $J_1$  interaction is “frustrated.”

angle or be straight, corresponding to the two types of tiles. Furthermore, using exact tensor contractions, we find that type I tiles form reflection-symmetry-broken sectors (either up or down triangles are ferromagnetic), whereas type II tiles (line straight across the hexagon) must appear in strings that form domain walls crossing the entire system between different reflection-symmetry-broken sectors. This is illustrated by a Monte Carlo sample in Fig. 8. In the ensemble of type I tiles the ferromagnetic up (down) triangles form effective degrees of freedom of a triangular Ising antiferromagnet. The residual entropy solely due to the type I tiles is thus  $\frac{1}{3}S_{\text{TLIAF}}$ . It would thus seem that the domain walls of type II tiles do not contribute to the residual entropy. To corroborate this we calculated the probability of finding a type II tile on some site using the contracted tensor network and found that this was 0 for all bond dimensions.

## VII. OUTLOOK

In this paper, we have introduced a general method to build contractible tensor networks for arbitrary frustrated Ising models. It relies on the identification of clusters in which the energy can be minimized independently and on a formulation of the partition function in terms of effective degrees of freedom that correspond to all the relevant ground states in each cluster. The construction is actually possible for any model with a discrete degree of freedom, for instance, Potts or clock models, and in any dimension.

To put this result in perspective, let us return to the core of the problem faced by tensor networks for frustrated systems, namely, the difficulty in numerically contracting tensors with simultaneously very large and very small elements at low temperatures because of their exponential dependence on the inverse temperature  $\beta$  with both positive and negative energies. This difficulty is very reminiscent of the sign problem in the quantum Monte Carlo, which excludes the investigation of the low-temperature properties of a quantum system if the off-diagonal matrix elements of the Hamiltonian are not all nonpositive, but which at the same time is basis dependent and can in principle be eliminated by a change of basis. What we have proposed here is a reformulation of the partition function in a basis where all elements are of the form  $e^{-\beta E}$ ,  $E \geq 0$ , leading to a tensor of relatively

modest dimension with only elements equal to 0 or 1 at zero temperature and to a contractible tensor with elements only in the interval  $[0,1]$  at any positive temperature. As for the sign problem in the quantum Monte Carlo the identification of the basis in general does not have a polynomial solution. Yet, we have shown that this is possible in practice, and we will give further examples in upcoming publications, where we will study finite-temperature properties and phase transitions in highly frustrated Ising systems. Study of the relation to, or combination with, the tropical tensor network approach to frustrated systems [40] would also be interesting and likely very fruitful.

Finally, we can consider the effect of quantum dynamics on the correlated phase spaces. Indeed we can write down PEPS wave functions by promoting the tiles to quantum degrees of freedom to effectively describe quantum corrections that would be present in any real-life material.

## ACKNOWLEDGMENTS

This work was initiated at the Centro de Ciencias de Benasque Pedro Pascual. F.V. thanks Michael Lawler for discussions about tensor networks and frustrated systems. This research was supported by ERC Grants QUTE (647905) and FWO (G0E1820N; L.V., B.V., F.V.) and the Swiss National Science Foundation (J.C., F.M.). The Monte Carlo computations were performed using the facilities of the Scientific IT and Application Support Center of the Ecole Polytechnique Fédérale de Lausanne.

## APPENDIX A: PSEUDOCODE

Algorithm 1 describes a method to find a simplex of maximal dimension that fits inside a convex set  $A$ . The idea is to take a small simplex inside  $A$  and try to make it bigger until it has the dimension of  $A$ . For this, the origin is first moved to the interior of the small simplex, and a vector orthogonal to the current simplex is constructed. One then looks for a vector in  $A$  of maximal overlap (in absolute value) with this vector. If the maximal overlap is 0, the simplex is of maximal dimension inside  $A$ ; if not, one adds the result to the simplex and starts over.

**Algorithm 1.** Build interior simplex of convex set  $A$ .

---

```

1:  $\vec{R} \leftarrow$  random vector
2:  $\vec{\alpha}_1 \leftarrow \max_{\vec{\alpha}} \vec{R} \cdot \vec{\alpha}$  with  $\alpha \in A$ 
3: While do
4:    $\vec{\beta} \leftarrow$  a point in simplex  $[\{\vec{\alpha}_i\}]$ 
5:   Translate  $\vec{\alpha}$ -space by  $-\vec{\beta}$ 
6:    $\{\vec{w}_i\} \leftarrow$  a basis of orthogonal vectors to all  $\{\vec{\alpha}_i\}$ 
7:   for  $\vec{v} \in \{\vec{w}_i, -\vec{w}_i\}$  do
8:      $\vec{\alpha} \leftarrow \max_{\vec{\alpha}} \vec{v} \cdot \vec{\alpha}$  with  $\alpha \in A$ 
9:     if  $\vec{v} \cdot \vec{\alpha} \neq 0$  then
10:       Add  $\vec{\alpha}$  to the set  $\{\vec{\alpha}_i\}$ 
11:       Return to the top of the while loop
12:   Stop the while loop
13: return  $\{\vec{\alpha}_i\}$ 

```

---

Algorithm 1 finds the interior simplex once the set  $A$  that solves the problem

$$E_u \leftarrow \max_{\vec{\alpha}} E, \quad \text{with} \quad \left\{ \begin{array}{l} H_u^{(\alpha)}(C) \geq E \quad \forall C \\ \sum_{c \in \mathcal{T}_u | n \in c} \alpha_n^c = 1 \end{array} \right\}, \quad (\text{A1})$$

has been found. For large clusters, where there is a large number of configurations, even finding this cluster can pose memory issues. Algorithm 2 offers a solution to this problem which automatically finds an interior simplex of Algorithm 1 using as few configurations  $C$  as possible. The algorithm turns the memory load into a time load.

The idea is to take a restricted set of configurations  $\{c_i\}$  and solve Eq. (A1) just for them, to estimate the full solution. This amounts to finding a convex set that contains  $A$ . We then look for an interior simplex of the solution set of this problem. Solving Eq. (A1), we get a temporary estimate for the ground-state energy. In each corner of the simplex, the  $\vec{\alpha}$  define a Hamiltonian associating an energy with all the configurations of the clusters. The estimate is compared to the energy of each configuration in each corner. If we find a configuration which has an energy below the estimate, the associated inequality is useful; the configuration is added to the set  $\{c_i\}$  and we restart. On the other hand, if in each corner, the energy of each configuration is greater than or equal to the estimate, we are sure to have a solution of Eq. (A1) with all configurations considered, and the problem is solved. This way, we work around most of the redundancy in the set of inequalities associated with all the configurations, and only inequalities that provide insight are used to build the convex set  $A$ .

**Algorithm 2.** Build interior simplex of  $A^u$ .

---



---

```

1:  $\{c_i\} \leftarrow$  choose some random configurations
2:
3: Add random configurations to  $\{c_i\}$  until there
4: is a finite  $E$  and a finite interior simplex
5:
6: While do
7:   for  $\vec{\alpha} \in \{\vec{\alpha}_i\}$  do
8:     for  $c : H_u^{\vec{\alpha}}(c) < E_{\text{temp}}$  do
9:       if  $c \notin \{c_i\}$  then
10:        Add  $c$  to  $\{c_i\}$ 
11:         $E_{\text{temp}} \leftarrow$  solve Eq. (A1) for configurations  $\{c_i\}$ 
12:         $\{\vec{\alpha}_i\} \leftarrow$  Update interior simplex for  $\{c_i\}$ 
13:        Return to the top of the while loop
14:   Stop the while loop
15: return  $E_{\text{temp}}, \{\vec{\alpha}_i\}$ 

```

---



---

Note that Algorithm 1 can just about handle the example from the paper, but larger clusters, say two or three kagome stars, could only be considered using Algorithm 2.

## APPENDIX B: ANALYSIS OF THE GROUND-STATE TILES

The ground-state tiles can be classified into two types: 48 type I tiles that have three nonoverlapping triangles of aligned

spins (i.e., ferromagnetic triangles) and 84 type II tiles that do not; see Fig. 10, where we have also drawn a line separating the up from the down spins. We proceed to understand the ground-state ensemble of this model by first characterizing the type I ensemble and then describing how type II tiles modify this picture.

### 1. Type I ensemble

The type I tiles are exactly all the configurations of  $u$  for which the three ferromagnetic triangles are never all pointing in the same direction. Each ferromagnetic triangle can be seen as an Ising degree of freedom; the type I tiles are thus all the configurations for which these three degrees of freedom are never aligned. The tiles can be separated into two subtypes by reflection symmetry: the tiles where the new Ising degrees of freedom live on up triangles and those where they live on down triangles. A global state made of tiling uniquely type I tiles can only be made of one of these subtypes, because the tiles in one subtype cannot be overlapping with the tiles in the other subtype. Therefore, the type I ensemble features reflection symmetry breaking.

The up (down) triangles are arranged as a triangular lattice, and we have seen that the type I tiles are *all* the configurations for which the three effective Ising degrees of freedom are not all aligned. So, there are no other constraints for tiling these type I tiles. It straightforwardly follows that the effective Ising degrees of freedom must act like the spins of an Ising antiferromagnet on the triangular lattice, a model whose residual entropy is known exactly [3]. The residual entropy of the type I ensemble is thus given by  $S = \frac{1}{3} S_{\text{TLIAF}}$ .

### 2. Type II ensemble

First, let us distinguish the type II tiles from the type I tiles. All the type II tiles have a line (indicating an interface between up and down spins) running *straight* across the hexagon. Conversely, all the configurations of the kagome star satisfying this description are type II tiles. In type I tiles, on the other hand, the lines separating up from down spins must only live on up (down) triangles. An immediate consequence of this is that type II tiles can connect type I tiles in different reflection symmetry sectors—if there are states containing the two types of tiles.

A key characteristic of a type II tile is the orientation of the line crossing the hexagon. We use this to identify three subtypes of type II tiles, illustrated in Fig. 12.

Making exact statements about how the tiles of the various subtypes can be matched together is not as easy as in the case of type I tiles. To see which subtypes can neighbor one another and how, we use a small tensor network construction and exact contractions. Imagine a patch of  $5 \times 5$  clusters where we restrict the center tile to one subtype of the type II tiles and ask what types the surrounding clusters may be, while satisfying the usual tiling rules. The tensor network for this is shown in Fig. 11. Note that this does not correspond to a finite system but, rather, a patch of  $5 \times 5$  in an infinite system. To each tensor neighboring the central cluster, we add a leg, allowing us to probe the local configuration. After contraction, the indices of the resulting tensor thus correspond to the labels

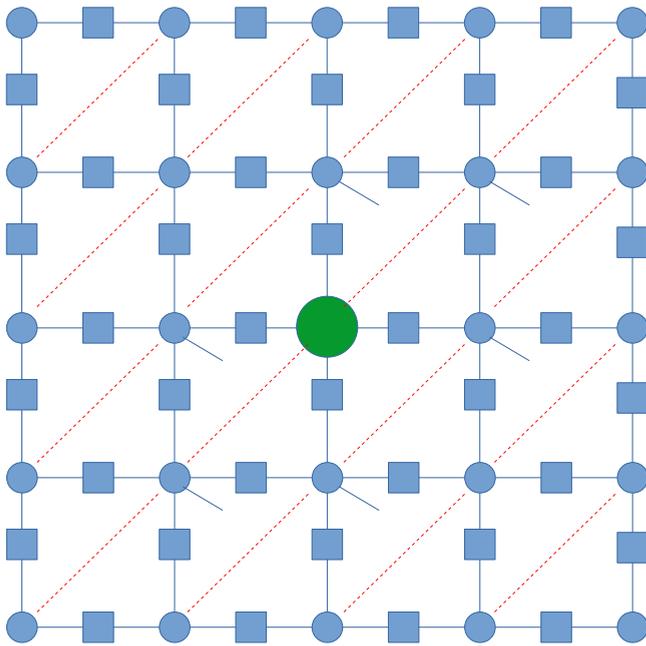


FIG. 11. The vertex tensors are  $\delta$  tensors representing the cluster configurations that make up the ground states; on the bonds are the usual  $P$  matrices enforcing tiling rules. Of course, in practice, we perform an SVD and exact truncation of these rank-deficient matrices to be able to perform computations more efficiently. The green vertex tensor in the middle is a  $\delta$  tensor that has been restricted to a single subtype of type II tiles. The tensor network has a square lattice shape, but in reality the clusters form a triangular lattice; the dotted red lines have therefore been added to indicate nearest neighbors. To the six  $\delta$  tensors describing the nearest neighbors of the central cluster we give an extra open index, allowing us to probe the local configuration.

of the tiles of the six nearest-neighbor clusters. If the value of the tensor at a certain set of indices is 0, it means that this configuration of neighboring clusters is not allowed. Some nonzero elements may become 0 if we consider a larger patch or even only if we consider the entire infinite plane; namely,

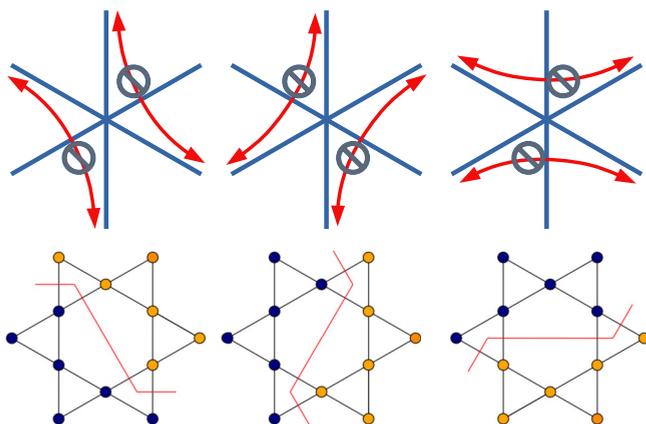


FIG. 12. Red arrows indicate which angles cannot be made by the different types of domain wall. These restrictions make it so that domain walls cannot make U-turns.

the configuration might be allowed locally but create some tiling issues at larger scales or at infinity.

The first result that we obtain is that a type II tile must have *exactly two* type II tiles of the same subtype as nearest neighbors. The type II tiles must thus make unending strings that conserve subtype, and these strings cannot cross or fuse. Additionally, we obtain from the forbidden local cluster configurations that these strings must either go straight or make  $120^\circ$  angles but cannot make sharp angles. Moreover, of those  $120^\circ$  angles, two of the six are forbidden (which two angles depends on the subtype), making it impossible for a string of type II tiles to close in on itself (under open boundary conditions). The forbidden angles with corresponding subtype are shown in Fig. 12.

We thus find that the type II tiles must form domain walls that extend the entire size of the system, separating different reflection-symmetry-broken sectors made up of type I tiles. A given domain wall can only have one specific symmetry-broken sector on either side, so two neighboring domain walls cannot be of the same type, but instead the types must alternate.

Finally, note that based on this analysis it is not clear whether the type II tiles are actually tessellable, and we only found an upper bound to their freedom. We do, however, have a whole lot of ground states from the Monte Carlo simulations, and so we know that this picture of domain walls is indeed correct.

### APPENDIX C: MONTE CARLO DETAILS

For the Monte Carlo simulations, as a complement to the standard single-spin-flip update which is rapidly failing, we use a dual-worm algorithm based on Ref. [12] as well as parallel tempering (also known as the temperature replica method). For this, we use 216 walkers, with a temperature associated with each walker. In a given Monte Carlo step, we first update each state with twice as many single-spin-flip attempts as there are sites in the system; we then perform worm updates until the total length of the worms corresponds to twice the number of dual sites in the system (see below); finally, we perform a parallel tempering step. At each step, detailed balance is respected.

In the dual-worm algorithm, the Ising model on a kagome lattice is first mapped onto a dimer model on the dice lattice according to

$$H = J_1 \sum_{\langle i,j \rangle} \sigma_i \sigma_j + J_2 \sum_{\langle\langle i,j \rangle\rangle} \sigma_i \sigma_j + J_3 \sum_{\langle\langle\langle i,j \rangle\rangle\rangle} \sigma_i \sigma_j \quad (C1)$$

$$= J_1 \sum_{\alpha} d_{\alpha} + J_2 \sum_{\Gamma_2} \prod_{\alpha \in \Gamma_2} d_{\alpha} + J_3 \sum_{\Gamma_3} \prod_{\alpha \in \Gamma_3} d_{\alpha}, \quad (C2)$$

where  $d_{\alpha} = \sigma_i \sigma_j$  if  $\alpha$  is the dual bond between the kagome lattice sites  $i$  and  $j$ , and where  $\Gamma_2$  ( $\Gamma_3$ ) goes through all the direct dimer paths connecting second-nearest-neighbor (third-nearest-neighbor) spins (Fig. 13). Building a loop update on this dimer model then corresponds to building a cluster in the original spin model. The loop is built respecting a local detailed balance condition such that, if the loop closes and the winding number of the loop in both directions on the torus is even, the update can be accepted. At the level of the

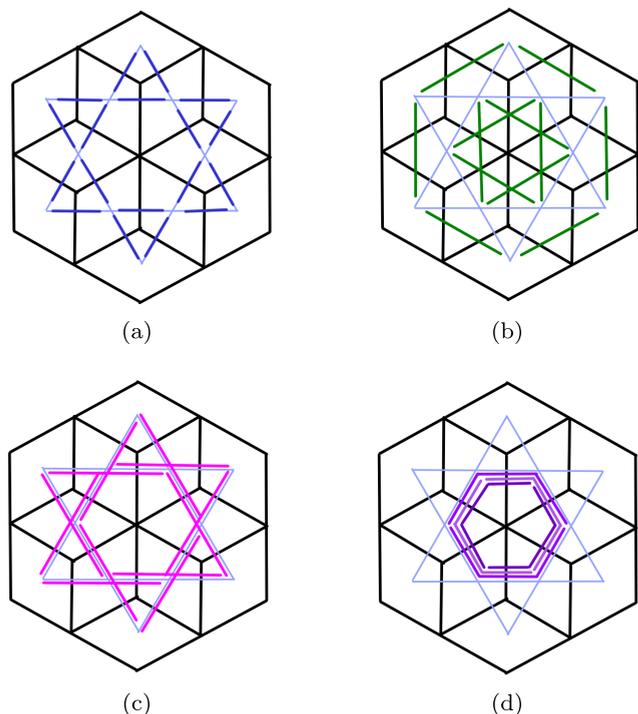


FIG. 13. Illustration of the further-neighbor interaction paths on a kagome lattice. (a) Nearest-neighbor interaction paths (coupling  $J_1$ ). (b) Next-nearest-neighbor interaction paths (coupling  $J_2$ ). (c) Third-nearest-neighbor interaction paths along the bonds of the kagome lattice (coupling  $J_3$ ). (d) Third-nearest-neighbor interaction paths, this time across the hexagons of the lattice (coupling  $J_3$ ).

dimer configuration, the local detailed balance is imposed by choosing the next direction to grow the loop uniformly at one step and based on a weight table at the next step (see Ref. [12] and references therein for the detailed balance proof and the weight table expressions for zero-bounce and one-bounce solutions). If the winding number is odd in either direction, the updated dimer configuration on the dice lattice does *not* map back to a periodic spin model on the kagome lattice. In such cases, additional loops are built until both winding numbers are even.

At the parallel tempering step, it is the detailed balance of the ensemble of walkers which is respected (see, for instance, [41] and references therein). This is done by going through pairs of configurations and accepting to swap with probability

$$A([\{\sigma_1\}, \beta_1), [\{\sigma_2\}, \beta_2]) \rightarrow [\beta_1 \leftrightarrow \beta_2]) = \min\{1, e^{(\beta_1 - \beta_2)(H\{\sigma_1\} - H\{\sigma_2\})}\}. \tag{C3}$$

At even Monte Carlo steps we go through pairs starting with even indexed temperatures, while at odd steps we go through pairs starting with odd-indexed temperatures.

These features of the Monte Carlo simulations allow one to reach the ground states. This is verified by computing the expectation value of the energy at the lowest temperature and checking that it is systematically within  $10^{-9}$  of the exact ground-state energies.

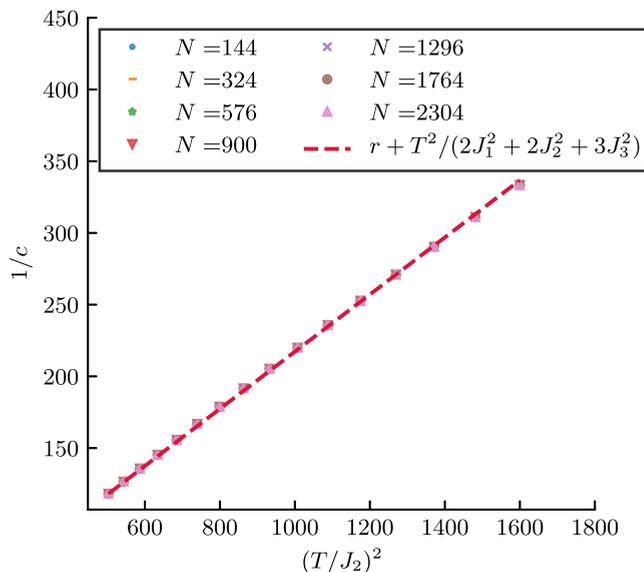


FIG. 14. Comparison of the specific heat near the maximal temperature, and the high-temperature expansion. The constant offset  $r \cong 18$  corresponds to a correction of order  $\beta^4$  in the high-temperature expansion, whose contribution to the entropy is negligible.

For each size, 16 independent runs are performed (for each run, 16 384 thermalization steps are followed by 1 048 576 measurement steps).

The specific heat per site  $c$  is computed using the variance of the energy (see, for instance, [42]) and the residual entropy per site is obtained from thermodynamic integration as

$$S = \ln 2 - \int_0^\infty \frac{c}{T} dT. \tag{C4}$$

Note that, in practice, we integrate numerically up to the maximal temperature  $T_{\max}/|J_2| = 40$ . For the temperatures from  $T_{\max}$  to  $\infty$ , one can compute the behavior of the specific heat from a high-temperature expansion. It is fairly easy to show

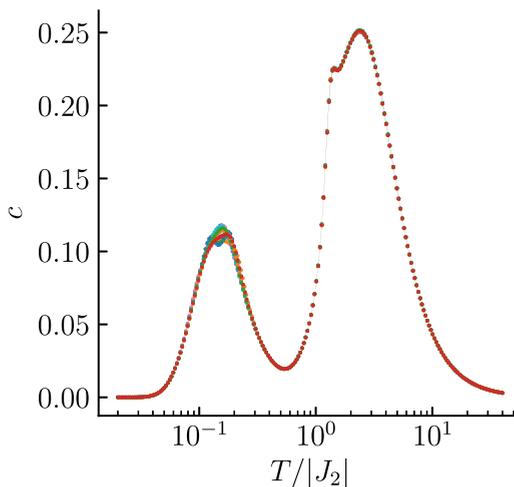


FIG. 15. Specific heat for independent runs. In the lower heap, the various runs do not agree (simulations for  $L = 12$ ).

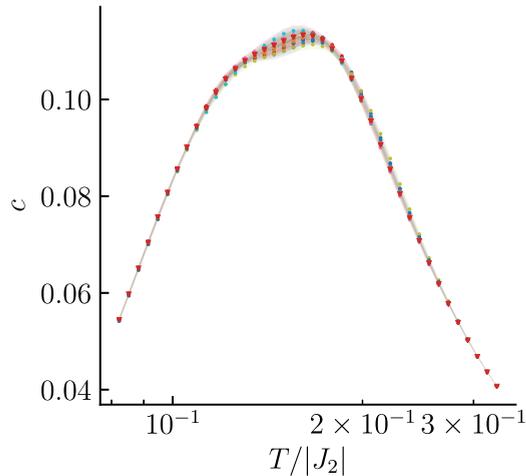


FIG. 16. Performing a bootstrap analysis, the various averages of the runs agree within the error bars (zooming in on the heap at a lower temperature for  $L = 12$ ).

that the first relevant term is

$$c_{\text{High } T} = \beta^2(2J_1^2 + 2J_2^2 + 3J_3^2) \quad (\text{C5})$$

and thus the integral  $\int_{T_{\text{max}}}^{\infty} \frac{c}{T} dT \cong 1.568 \times 10^{-4}$  [plotting  $1/c$  as a function of  $T^2$ , we can check that at  $T \cong T_{\text{max}}$  this first term already captures the behavior very well (Fig. 14)].

The specific heat in the lowest heap shows a dependence on the run for large sizes (Fig. 15). This is compensated for by taking the average over the 16 simulations. By a bootstrap

analysis, we show that the error bars obtained from merging the 16 independent simulations are reasonable (Fig. 16). The error bars on the specific heat (2 standard deviations) are used to give the error bars on the residual entropy by integrating the smallest (largest) possible value of the specific heat over  $T$  at any temperature. Finally, we note that the residual entropy can alternatively be computed by integrating over the energy (see, e.g., Ref. [43] and references therein). We did this and kept only those sizes for which the simulations had been run long enough that the two ways of computing the residual entropy would agree within the error bars.

This model seems to have extremely strong, hard-to-characterize, finite-size effects. We show the residual entropy as a function of the inverse of  $N$ , the number of spins, and  $L$ , the linear system size, in Fig. 9. Justifying finite-size corrections is a challenge, often requiring a preexisting understanding of the ground-state phase, and here we only show these two graphs as a guide for the eye. The extrapolation as a function of  $1/N$  would seem to work best, at least compared to the tensor network result. However, the slope of  $S(N)$  indicates a huge prefactor to the number of ground states that we cannot explain. The extrapolation in  $1/L$  would seem most plausible based on the observation that the type II tiles form domain walls, especially since they are irrelevant to the extensive entropy. But the extrapolation in  $1/L$  does not look too convincing and would dramatically underestimate the lower bound of  $S = \frac{1}{3}S_{\text{TLIAF}}$ . To solve this problem, one would need to study even larger system sizes, which turns out to be very difficult with the Monte Carlo, at least with our algorithm.

- 
- [1] C. Lacroix, P. Mendels, and F. Mila, *Introduction to Frustrated Magnetism: Materials, Experiments, Theory*, Vol. 164 (Springer Science and Business Media, New York, 2011).
  - [2] C. L. Henley, The ‘‘Coulomb phase’’ in frustrated systems, *Annu. Rev. Condens. Matter Phys.* **1**, 179 (2010).
  - [3] G. H. Wannier, Antiferromagnetism. The triangular Ising net, *Phys. Rev.* **79**, 357 (1950).
  - [4] K. Kanô and S. Naya, Antiferromagnetism. The Kagomé Ising Net, *Prog. Theor. Phys.* **10**, 158 (1953).
  - [5] P. W. Kasteleyn, Dimer statistics and phase transitions, *J. Math. Phys.* **4**, 287 (1963).
  - [6] T. D. Schultz, D. C. Mattis, and E. H. Lieb, Two-dimensional Ising model as a soluble problem of many fermions, *Rev. Mod. Phys.* **36**, 856 (1964).
  - [7] E. H. Lieb, Residual entropy of square ice, *Phys. Rev.* **162**, 162 (1967).
  - [8] R. J. Baxter, *Exactly Solved Models in Statistical Mechanics* (Academic Press, New York, 1982).
  - [9] R. H. Swendsen and J.-S. Wang, Nonuniversal Critical Dynamics in Monte Carlo Simulations, *Phys. Rev. Lett.* **58**, 86 (1987).
  - [10] U. Wolff, Collective Monte Carlo Updating for Spin Systems, *Phys. Rev. Lett.* **62**, 361 (1989).
  - [11] Y. Wang, H. De Sterck, and R. G. Melko, Generalized Monte Carlo loop algorithm for two-dimensional frustrated Ising models, *Phys. Rev. E* **85**, 036704 (2012).
  - [12] G. Rakala and K. Damle, Cluster algorithms for frustrated two-dimensional Ising antiferromagnets via dual worm constructions, *Phys. Rev. E* **96**, 023304 (2017).
  - [13] 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).
  - [14] L. Vanderstraeten, B. Vanhecke, and F. Verstraete, Residual entropies for three-dimensional frustrated spin systems with tensor networks, *Phys. Rev. E* **98**, 042145 (2018).
  - [15] 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).
  - [16] J. Haegeman and F. Verstraete, Diagonalizing transfer matrices and matrix product operators: A medley of exact and computational methods, *Annu. Rev. Condens. Matter Phys.* **8**, 355 (2017).
  - [17] W. Huang, D. A. Kitchaev, S. T. Dacek, Z. Rong, A. Urban, S. Cao, C. Luo, and G. Ceder, Finding and proving the exact ground state of a generalized Ising model by convex optimization and MAX-SAT, *Phys. Rev. B* **94**, 134424 (2016).
  - [18] The tensors are the same for the triangular lattice Ising antiferromagnet.
  - [19] Indeed, there cannot be a ground state containing a ferromagnetic down triangle: This would mean that one of the triangular Hamiltonians is not minimized.

- [20] In programmatic terms, the Hamiltonian tessellation can be thought of as a class of Hamiltonian splittings described by the choice of the cluster  $u$  and where choosing the values for the weights  $\{\alpha_n^c\}$  defines an instance of the class.
- [21] V. Zauner-Stauber, L. Vanderstraeten, M. T. Fishman, F. Verstraete, and J. Haegeman, Variational optimization algorithms for uniform matrix product states, *Phys. Rev. B* **97**, 045145 (2018).
- [22] M. T. Fishman, L. Vanderstraeten, V. Zauner-Stauber, J. Haegeman, and F. Verstraete, Faster methods for contracting infinite two-dimensional tensor networks, *Phys. Rev. B* **98**, 235148 (2018).
- [23] L. Vanderstraeten, J. Haegeman, and F. Verstraete, Tangent-space methods for uniform matrix product states, *SciPost Phys. Lect. Notes* **7**, 1 (2019).
- [24] T. Nishino and K. Okunishi, Corner transfer matrix renormalization group method, *J. Phys. Soc. Jpn.* **65**, 891 (1996).
- [25] Z. Zhu and H. G. Katzgraber, Do tensor renormalization group methods work for frustrated spin systems? [arXiv:1903.07721v1](https://arxiv.org/abs/1903.07721v1).
- [26] M. Kaburagi and J. Kanamori, A method of determining the ground state of the extended-range classical lattice gas model, *Prog. Theor. Phys. Kyoto* **54**, 30 (1975).
- [27] T. Kudo and S. Katsura, A method of determining the orderings of the Ising model with several neighbor interactions under the magnetic field and applications to hexagonal lattices, *Prog. Theor. Phys.* **56**, 435 (1976).
- [28] The method described in Ref. [17] is dual to their basic polytope method, which is itself similar to Kanamori's method developed in Ref. [26]. The main differences are that Kanamori's method has weaker convergence properties but gives lower bounds for all the couplings at once.
- [29] M. Grötschel, L. Lovász, and A. Schrijver, *Geometric Algorithms and Combinatorial Optimization* (Springer-Verlag, Berlin, 1993).
- [30] A simplex is the generalization to a high dimension of triangles and tetrahedrons.
- [31] H. Wang, Proving theorems by pattern recognition—II, *Bell Syst. Tech. J.* **40**, 1 (1961).
- [32] R. M. Robinson, Undecidability and nonperiodicity for tilings of the plane, *Invent. Mathematicae* **12**, 177 (1971).
- [33] H. Wang, Notes on a class of tiling problems, *Fund. Mathematicae* **82**, 295 (1975).
- [34] Z. Wang and M. Navascués, Two-dimensional translation-invariant probability distributions: Approximations, characterizations and no-go theorems, *Proc. R. Soc. A* **474**, 20170822 (2018).
- [35] T. Takagi and M. Mekata, Magnetic ordering of Ising spins on kagomé lattice with the 1st and the 2nd neighbor interactions, *J. Phys. Soc. Jpn.* **62**, 3943 (1993).
- [36] M. Wolf and K. D. Schotte, Ising model with competing next-nearest-neighbour interactions on the kagome lattice, *J. Phys. A: Math. Gen.* **21**, 2195 (1988).
- [37] T. Mizoguchi, L. D. C. Jaubert, and M. Udagawa, Clustering of Topological Charges in a Kagome Classical Spin Liquid, *Phys. Rev. Lett.* **119**, 077207 (2017).
- [38] I. A. Chioar, N. Rougemaille, and B. Canals, Ground-state candidate for the classical dipolar kagome Ising antiferromagnet, *Phys. Rev. B* **93**, 214410 (2016).
- [39] J. Hamp, R. Moessner, and C. Castelnovo, Supercooling and fragile glassiness in a dipolar kagome Ising magnet, *Phys. Rev. B* **98**, 144439 (2018).
- [40] J.-G. Liu, L. Wang, and P. Zhang, Tropical tensor network for ground states of spin glasses, [arXiv:2008.06888](https://arxiv.org/abs/2008.06888).
- [41] H. G. Katzgraber, S. Trebst, D. A. Huse, and M. Troyer, Feedback-optimized parallel tempering Monte Carlo, *J. Stat. Mech.: Theory Exp.* (2006) P03018.
- [42] A. W. Sandvik, Computational studies of quantum spin systems, in *Lectures on the Physics of Strongly Correlated Systems XIV: Fourteenth Training Course in the Physics of Strongly Correlated Systems*, edited by A. Avella and F. Mancini, AIP Conf. Proc. No. 1297 (AIP, New York, 2010).
- [43] M. Žukovič, Residual entropy of spin-s triangular Ising antiferromagnet, *Eur. Phys. J. B* **86**, 283 (2013).