


Pyrochlore $S = \frac{1}{2}$ Heisenberg antiferromagnet at finite temperature

Robin Schäfer^{1,*}, Imre Hagymási^{1,2,†}, Roderich Moessner^{1,‡} and David J. Luitz^{1,§}

¹Max Planck Institute for the Physics of Complex Systems, Noethnitzer Strasse 38, 01187 Dresden, Germany

²Strongly Correlated Systems “Lendület” Research Group, Institute for Solid State Physics and Optics, Wigner Research Centre for Physics, Budapest H-1525 P.O. Box 49, Hungary

 (Received 16 April 2020; revised 5 June 2020; accepted 15 July 2020; published 4 August 2020)

We use a combination of three computational methods to investigate the notoriously difficult frustrated three-dimensional pyrochlore $S = \frac{1}{2}$ quantum antiferromagnet, at finite temperature T : canonical typicality for a finite cluster of $2 \times 2 \times 2$ unit cells (i.e., 32 sites), a finite- T matrix product state method on a larger cluster with 48 sites, and the numerical linked cluster expansion (NLCE) using clusters up to 25 lattice sites, including nontrivial hexagonal and octagonal loops. We calculate thermodynamic properties (energy, specific heat capacity, entropy, susceptibility, magnetization) and the static structure factor. We find a pronounced maximum in the specific heat at $T = 0.57J$, which is stable across finite size clusters and converged in the series expansion. At $T \approx 0.25J$ (the limit of convergence of our method), the residual entropy per spin is $0.47k_B \ln 2$, which is relatively large compared to other frustrated models at this temperature. We also observe a nonmonotonic dependence on T of the magnetization at low magnetic fields, reflecting the dominantly nonmagnetic character of the low-energy states. A detailed comparison of our results to measurements for the $S = 1$ material $\text{NaCaNi}_2\text{F}_7$ yields a rough agreement of the functional form of the specific heat maximum, which in turn differs from the sharper maximum of the heat capacity of the spin ice material $\text{Dy}_2\text{Ti}_2\text{O}_7$.

DOI: [10.1103/PhysRevB.102.054408](https://doi.org/10.1103/PhysRevB.102.054408)

I. INTRODUCTION

The pyrochlore lattice, composed of corner-sharing tetrahedra, is a common motif in materials chemistry; in the context of magnetic materials, it has been prominent in a range of rare-earth [1] and spinel compounds [2,3]. Pyrochlore magnets and models have played a tremendously important role in the history of frustrated magnetism and topological condensed matter physics. One of the foundational publications, in 1956, was Anderson’s identification of the classical pyrochlore Ising magnet [4] as an interesting model system. Now called spin ice [5], this is a topological magnet exhibiting an emergent gauge field and fractionalized excitations [6].

The classical Heisenberg model on the pyrochlore lattice, following a pioneering study by Villain [9], turned out to be the first classical Heisenberg spin liquid [10]. This undergoes a very delicate order-by-disorder transition for large spins, as the zero-point energy induced by quantum fluctuations favours a subset of collinear states [11,12]. Beyond this

(semi)classical limit of large spin $S \rightarrow \infty$, little is known reliably about the properties of the pyrochlore quantum Heisenberg model.

This is because the properties of the lattice conspire to frustrate not only magnetic order, but also attempts to apply standard theoretical and numerical approaches. The presence of a macroscopic number (“flat band”) of gapless excitations in most bare models precludes standard perturbative schemes and mean-field theories. This reflects the fact that fluctuations are typically very strong, the basic ingredient via which frustrated magnets avoid ordering. For this reason, even the relatively “high” dimensionality $d = 3$, often considered almost homologous with proximity to mean-field behavior, is a hindrance rather than a help: the most unbiased method, exact diagonalization, breaks down already for small linear system sizes L , as the Hilbert space dimension grows exponentially with L^3 . Similarly, DMRG-based methods are well known to struggle beyond $d = 1$, while geometric frustration yields a sign problem in quantum Monte Carlo. In this sense, the pyrochlore magnet is even less tractable than the notoriously enigmatic kagome $S = \frac{1}{2}$ Heisenberg magnet, for which decades of intensive interest have not yielded a consensus on the nature of its ground state. For the pyrochlore lattice, even a reliable ground-state energy estimate is lacking, and the proposed ground states depend strongly on the method used to study it [13–24].

The pyrochlore $S = \frac{1}{2}$ Heisenberg magnet thus has many ingredients that make it one of the most likely candidates for realizing new and exotic phases of matter, *inter alia*, quantum spin liquid states: it harbors both promise and obstacles, albeit in a somewhat imbalanced way.

*schaefer@pks.mpg.de

†hagymasi@pks.mpg.de

‡moessner@pks.mpg.de

§dluitz@pks.mpg.de




FIG. 1. Heat capacity per spin in different pyrochlore magnets, for a detailed account see Sec. IV. The red curve represents the converged part of our results from the NLCE for the $S = \frac{1}{2}$ Heisenberg model. Results for a single tetrahedron for $S = \frac{1}{2}$ (dashed) for $S = \infty$ (classical case, dots) have T scaled to match in the high- T limit. Symbols are for experiments on $\text{NaCaNi}_2\text{F}_7$ [7], a $S = 1$ (approximate) Heisenberg magnet, and the Ising spin ice $\text{Dy}_2\text{Ti}_2\text{O}_7$ [8], both scaled in T such that their maxima coincide with that of the $S = \frac{1}{2}$ model (with a factor $\ln 2 / \ln 3$ to account for the larger $S = 1$ entropy). Inset similarly shows entropy per spin. We use natural units in the paper $k_B = 1$, $\hbar = 1$, $J = 1$.

In view of all the above-mentioned difficulties, the best source for information on these low temperature phases are experiments. In particular, since scattering neutrons off a frustrated magnet is *a priori* no more involved than off an unfrustrated one, $d = 3$ permits the straightforward study of bulk samples. On the other hand, experiments bring their own challenges, since it is often difficult to determine the best low energy theory and to identify materials which are close to theoretical models.

Up to now several compounds are known whose magnetic properties can be described by Heisenberg-like models on the pyrochlore lattice [25–27]. While a good realization of an $S = 1$ Heisenberg magnet is now available [7,28,29], an isotropic nearest-neighbor-dominated $S = \frac{1}{2}$ material still remains on the wish list.

As experiments necessarily involve nonzero temperatures, theory is compelled to study this setting. Our study is therefore devoted to the $S = \frac{1}{2}$ pyrochlore magnet at finite temperature. We focus on the thermodynamics—susceptibility and in particular specific heat, Fig. 1, and we also consider the spin correlations in the form of the momentum resolved static structure factor.

The remainder of this account is structured as follows. In Sec. IA we provide a summary of our results. Section II introduces the model and observables. The bulk of the technical advances are bundled into Sec. III, with material on canonical typicality, the numerical linked cluster expansion, exact diagonalization, as well as DMRG. This may be skipped on first reading (as well as by the reader not interested in the underlying methodology). Our results on thermodynamics (specific heat, susceptibility in zero and nonzero fields) as well as spin correlators are presented in Sec. IV. We close with a broader discussion and an outlook in Sec. V.

A. Summary of results

1. Methods

Our quest to make progress has a considerable purely technical component. This involves efforts along two, principally computational, axes.

First, we devise a high-order numerical linked cluster expansion (NLCE) [30–32] for the pyrochlore lattice. This approach has been used before, and our contribution is to push the expansion—based on tetrahedral clusters which were shown to be well suited for the corner-sharing tetrahedra of the pyrochlore lattice [33–38]—to significantly higher orders. We reach clusters of up to eight tetrahedra, involving the full solution of clusters of up to 25 spins. Previous works considered mostly NLCE up to $n = 5$ tetrahedra, corresponding to clusters of up to 16 spins. Going to such high order allows a significantly improved exploration of the low temperature regime, in particular permitting a more extended and controlled use of Euler transforms to extrapolate the results to low temperature.

Indeed, high expansion orders are essential to capture a range of physical processes. Concretely, up to the sixth nearest-neighbor hop, the pyrochlore (and, indeed, the kagome lattice) is equivalent to a Husimi cactus [a Cayley tree of tetrahedra (triangles)], and many series expansions are “trivial” to high orders, e.g., with degeneracy lifting only occurring at *eighth* order in perturbation theory in a high-temperature kagome [39] or a strong-coupling pyrochlore expansion [40]. In a similar vein, the importance of resonance processes on more extended clusters is a recurring theme in the study of frustrated magnets. The clusters included in our expansion host not only the simplest hexagonal loop motifs but also loops of eight spins and longer decorated versions thereof (the longest loop consisting of eight tetrahedra cf. Fig. 14), which crucially encode the three-dimensional structure of the lattice.

The second technical axis involves a finite-temperature DMRG analysis of the pyrochlore $S = \frac{1}{2}$ magnet using finite clusters. Our results demonstrate that finite-temperature DMRG is a powerful approach, feasible even in this challenging three-dimensional setup down to nontrivial temperatures. Here we use a “snake” path through the lattice to map the system to one dimension with long-range interactions. We take advantage of the $\text{SU}(2)$ symmetry of the model and keep $\text{SU}(2)$ block states up to $\chi = 10\,000$ [$\sim 40\,000U(1)$ equivalent] and consider clusters up to 48 sites with periodic boundaries.

Taken together, these approaches permit us to reach converged results at temperatures down to around $T = 0.25$ [where the exchange constant of the Heisenberg model has been set to unity, Eq. (1)] for the NLCE, and down to $T = 0.6$ for DMRG.

2. Observables

In the zero-field specific heat we resolve a pronounced maximum at $T = 0.57$. Crucially, at $T = 0.25$, where the specific heat has dropped well below its maximum value, the residual entropy is still around 0.33, i.e., 47% of the value of a free spin of $\ln 2$. This demonstrates the persistence of the spectral weight downshift characteristic of frustrated magnets

to this case. We discuss implications of this observation in detail, in particular in comparison with the kagome magnet, as well as two simple tetrahedral models, on top of the experimental results on two pyrochlore magnets: the $S = 1$ Heisenberg antiferromagnet $\text{NaCaNi}_2\text{F}_7$ [7], and the classic Ising spin ice $\text{Dy}_2\text{Ti}_2\text{O}_7$ [8]. We find that our model at $T = 0.25$, Fig. 1, *has a higher low- T entropy than all of these.*

This entropy at $T = 0.25$ is in particular much greater than that proposed for singlet subspaces in resonating valence bond type effective theories, where pairs of spins at the endpoints of a bond form a singlet bond, with resonance processes supplying quantum dynamics between these valence bond configurations [13,16,17,41]. Indeed, in this regime there is considerable admixture of triplet components in each tetrahedral wave function, reflecting the inability of neighboring tetrahedra to be in singlet states simultaneously. In the magnetization curves this is reflected in a nonmonotonic temperature dependence for fixed intermediate fields: upon cooling from the maximum magnetization, the entropy of the magnetic excitations loses out to the singlet-dominated low-energy sector; while at high temperatures, the magnetization assumes a conventional asymptotic $1/T$ behavior. The maximum disappears at zero field, where there is no magnetization in the absence of time-reversal symmetry breaking; and at high fields, where a conventional monotonic paramagnetic magnetization curve is found.

For the magnetic-field dependence of the specific heat, we find a continuous drift of the location of its maximum to higher temperatures; at the same time, the amplitude of the maximum changes nonmonotonically, first decreasing and then increasing again.

The spin correlators in turn, exhibit the by now familiar structure of incipient bowties, commonly found in various magnets on the pyrochlore lattice [7,13,15,24,42,43]. These reflect the emergent gauge field and while they become arbitrarily sharp in the cases of classical magnets, their finite width indicates the presence of a nonzero net moment on the tetrahedra, on account of the above-mentioned inability to have tetrahedra sharing a spin to be in a singlet configuration simultaneously.

II. MODEL AND OBSERVABLES

We focus on the isotropic spin $S = \frac{1}{2}$ Heisenberg antiferromagnet

$$H = \sum_{\langle i,j \rangle} \vec{S}_i \cdot \vec{S}_j + h \sum_i S_i^z. \quad (1)$$

The spins reside on the sites i of the three-dimensional (3D) pyrochlore lattice, which is a face centered cubic lattice with lattice vectors $\vec{a}_1 = \frac{1}{2}(1, 1, 0)^T$, $\vec{a}_2 = \frac{1}{2}(1, 0, 1)^T$, $\vec{a}_3 = \frac{1}{2}(0, 1, 1)^T$ and a tetrahedral basis given by $\vec{b}_0 = (0, 0, 0)^T$, $\vec{b}_1 = \frac{1}{4}(1, 1, 0)^T$, $\vec{b}_2 = \frac{1}{4}(1, 0, 1)^T$, $\vec{b}_3 = \frac{1}{4}(0, 1, 1)^T$, such that each lattice point can be expressed by

$$\vec{R}_{\alpha, n_1, n_2, n_3} = n_1 \vec{a}_1 + n_2 \vec{a}_2 + n_3 \vec{a}_3 + \vec{b}_\alpha, \quad (2)$$

with integer n_1, n_2, n_3 and $\alpha \in \{0, 1, 2, 3\}$.

The sum $\langle i, j \rangle$ in Eq. (1) runs over nearest-neighbor bonds of the pyrochlore lattice. In the absence of a magnetic field $h = 0$, the model is $\text{SU}(2)$ symmetric.

In this work we focus on thermodynamic observables: the heat capacity at fixed volume C_V , the magnetic susceptibility χ , and the entropy S . We also consider the static spin structure factor $S(\vec{Q})$. In the following definitions we use the canonical ensemble averages $\langle \bullet \rangle_\beta = \frac{1}{Z} \text{Tr}(e^{-\beta H} \bullet)$. As usual, $\beta = 1/T$ denotes the inverse temperature, and $Z = \text{Tr}(e^{-\beta H})$ is the partition function.

The heat capacity is obtained either from the temperature derivative of the internal energy $\langle H \rangle_\beta$ or from the fluctuations of the energy:

$$C_V = \frac{\partial \langle H \rangle_\beta}{\partial T} = \beta^2 (\langle H^2 \rangle_\beta - \langle H \rangle_\beta^2). \quad (3)$$

Similarly, we obtain the magnetic susceptibility χ , defined by the change of the magnetization in z direction with respect to a change of the field h in z direction from the fluctuations of the magnetization:

$$\chi = \frac{\partial \langle m_z \rangle_\beta}{\partial h} = \beta (\langle m_z^2 \rangle_\beta - \langle m_z \rangle_\beta^2), \quad (4)$$

with $m_z = \sum_i S_i^z$.

In the $\text{SU}(2)$ -symmetric case, the susceptibility can be also expressed as

$$\chi = \frac{\beta}{3N} \sum_{ij} \langle \vec{S}_i \cdot \vec{S}_j \rangle_\beta. \quad (5)$$

The thermodynamic entropy S can be calculated using the definition of the free energy:

$$S = \ln Z + \beta \langle H \rangle_\beta. \quad (6)$$

The static structure factor can be obtained from the Fourier transformation of the spin-spin correlations (the factor $4/3$ stems from the normalization $1/[S(S+1)]$ for spin $S = \frac{1}{2}$):

$$S(\vec{Q}) = \frac{4}{3N} \sum_{ij} \langle \vec{S}_i \cdot \vec{S}_j \rangle_\beta \cos [\vec{Q} \cdot (\vec{R}_i - \vec{R}_j)], \quad (7)$$

where \vec{R}_i denote the real-space coordinates of sites according to Eq. (2).

III. METHODS

This section is devoted to a detailed exposition of the methods used to obtain the results presented in the following section. It can safely be skipped at first reading, as well as by the reader primarily interested in the behavior of the observables, rather than details on how they were obtained.

A. Canonical typicality

Calculating thermodynamic expectation values is possible via the density matrix $\rho_\beta = \frac{1}{Z} e^{-\beta H}$, which can be calculated from all eigenvalues and eigenvectors of the Hamiltonian H . Due to the exponential scaling of the Hilbert space dimension with system size, this is impractical for systems with more

than $\gtrsim 25$ spins (we discuss how to perform full diagonalization for such systems in Sec. III D). The concept of quantum typicality [44,45] permits a different approach, which has its foundation in Lévy's lemma. It can be summarized in the statement that for the vast majority of wave functions $|\psi\rangle$, the state $|\beta\rangle = e^{-\beta/2H}|\psi\rangle$ is typical [46–48] for the canonical ensemble.

In practice, this means that, starting from a random wave function $|\psi\rangle$, one can calculate finite temperature expectation values of observables O by

$$\langle O \rangle_\beta = \frac{\text{Tr}(e^{-\beta H} O)}{Z} = \frac{\langle \beta | O | \beta \rangle}{\langle \beta | \beta \rangle} + \mathcal{O}(e^{-N}). \quad (8)$$

The statistical error of this replacement is exponentially small in system size [49] and can be estimated (and reduced) by sampling over random (infinite temperature) initial wave functions $|\psi\rangle$. The application of $e^{-(\beta/2)H}$ corresponds to imaginary time evolution up to $\beta/2$ (the factor 1/2 stems from a symmetric splitting of the exponential) and can be carried out efficiently using Krylov space techniques [48,50–53], which is commonly known as the finite temperature Lanczos method [52]. The main advantage of this technique is that it can be carried out storing only the vectors of the Krylov space spanned by $|\psi\rangle$, while the application of the Hamiltonian to vectors during the Lanczos algorithm is either performed using a sparse matrix representation, or an on the fly generation of the corresponding matrix elements, making very large system sizes accessible which are comparable to Lanczos ground-state calculations.

We note that the same techniques for reducing the Hamiltonian to its symmetry sectors discussed in Sec. III D can be readily applied here.

B. Finite temperature calculations with matrix product states

Matrix-product-state (MPS) [54,55] based algorithms also provide a way to address the equilibrium thermodynamics of many-body quantum systems [56–59]. One of the most widely used methods is the purification of the finite-temperature density matrix [56,59,60]. The idea behind this approach is that one can interpret the density matrix ρ_P as a partial trace of a Schmidt decomposition of a pure state $|\Psi\rangle$ in an enlarged Hilbert space:

$$|\Psi\rangle = \sum_\alpha s_\alpha |\alpha\rangle_P |\alpha\rangle_A \rightarrow \rho_P = \text{Tr}_A |\Psi\rangle \langle \Psi|, \quad (9)$$

where P and A denote the physical and auxiliary system, respectively. The state $|\Psi\rangle$ can be easily constructed by simply creating a copy (auxiliary system) of the physical system and generating maximally entangled bonds between each physical site and its auxiliary site. This can be achieved by creating an entangler Hamiltonian whose ground state corresponds to the maximally entangled initial state. In our case this Hamiltonian simply reads

$$H_{\text{entangler}} = \sum_i \vec{S}_i \cdot \vec{S}_{a(i)}, \quad (10)$$

where $a(i)$ denotes the auxiliary site belonging to site i and the sum is performed over the physical sites. One can easily see that the density matrix calculated from this ground

state $|\Psi_{\beta=0}\rangle$ corresponds to infinite temperature. Any finite-temperature density matrix can be obtained by performing an imaginary time evolution on the physical system and tracing out the auxiliary degrees of freedom:

$$|\Psi_\beta\rangle = e^{-\beta H/2} |\Psi_{\beta=0}\rangle \rightarrow \rho_\beta = \text{Tr}_A |\Psi_\beta\rangle \langle \Psi_\beta|. \quad (11)$$

As a matter of fact, any expectation $\langle O \rangle_\beta$ can be directly evaluated from $|\Psi_\beta\rangle$:

$$\langle O \rangle_\beta = \frac{\langle \Psi_\beta | O | \Psi_\beta \rangle}{\langle \Psi_\beta | \Psi_\beta \rangle}. \quad (12)$$

This provides a great advantage of this method, since thermodynamic quantities can directly be obtained by simulating the density matrix, rather than averaging over low entanglement pure states [57,58], and therefore results are free of statistical errors. By design, this method is most efficient in one dimension. In order to use it for a three-dimensional system, one has to put a “snake” path through the lattice sites to map the original problem to a one-dimensional equivalent one, which contains long-range couplings between the lattice sites. This is the main difficulty of this approach, since the MPSs need to encode a large amount of entanglement, i.e., if we think of the area law for a moment (valid only for ground states), the bond dimension should scale exponentially in ℓ^2 (ℓ is the linear size of the 3D system) to accurately represent the many-body state. This obviously limits the feasible system sizes and the accessible temperatures. The presence of long-range couplings in the one-dimensional topology poses another difficulty regarding the choice of the time evolution method [61]. The time-evolving block decimation (TEBD) [62,63] is very effective if the couplings are short ranged, otherwise one has to subsequently apply a series of swap gates to move distant sites next to each other so that the time evolving operator can be applied. This swapping procedure is extremely slow and becomes very inefficient as the bond dimension is increased. The Krylov method [64] is capable of handling long-range interactions by default, but already in the early stages of the imaginary-time evolution the Krylov vectors become strongly entangled, making the calculation unfeasible. To reach physically relevant temperatures, we demonstrate that the time-dependent variational principle (TDVP) [65,66] provides an effective way. Although it introduces another source of error by projecting the evolution vector onto the MPS manifold, this error is usually much smaller than the truncation error. At this point we also have to mention that it is not straightforward to evolve $|\Psi_{\beta=0}\rangle$ directly with TDVP [61]. Since the initial state is essentially a product state a naive evolution with TDVP would be simply wrong due to the loss of long-range interactions already in the first projection step. To overcome this difficulty we apply the same trick that has been successfully applied in real-time evolution [61], namely, we generate an initial state with an artificially enlarged bond dimension. This is achieved by finding the ground state of the Hamiltonian:

$$H_{\text{DMRG}} = H_{\text{entangler}} + aH(h=0), \quad (13)$$

where the parameter a is being varied. We start with $a = 1$ and perform 20 sweeps, then we reduce it by a factor of 10 each time. During the sweeps a single-site version [67] of the density-matrix renormalization group (DMRG) algorithm

[55,68–70] is applied with subspace expansion and setting the truncation error to zero. Five stages are performed altogether. In the last stage we set $a = 0$ and perform three additional sweeps. We demonstrate that this procedure removes the above bottleneck of TDVP also for imaginary-time evolution. In addition, to encode the large amount of entanglement, the compression of the many-body states must be very efficient. To this end we exploit the $SU(2)$ symmetry of the model and keep block states up to $\sim 10\,000$ [$\sim 40\,000$ $U(1)$ equivalent] to minimize the truncation error as much as possible [71,72]. It is worth mentioning that higher bond dimensions can be achieved in a ground-state search [73], where one can take advantage of single-site DMRG as well as parallelizing the computation in real space to reduce memory usage and computation time, respectively. In our case, however, the two-site TDVP update scheme needs to be used and the serial solution of the TDVP equations is crucial.

C. Numerical linked cluster expansion

For studying three-dimensional frustrated quantum magnets most controlled algorithms are restricted to a small number of spins and thus have a hard time capturing the three-dimensional structure and its correlations. Using a systematic high temperature series expansion opens up the possibility to obtain reliable results in the thermodynamic limit. The numerical linked cluster expansion (NLCE) is able to determine any extensive property P in the high temperature regime. NLCE has been applied to various geometries like the square lattice, kagome lattice, or pyrochlore lattice and has provided new insights into these systems [30–32,34,74,75] such as the transition of different phases in quantum systems [36,38,76] or a deeper understanding of real materials [33,35,37]. Moreover, the generality allows the application of this algorithm to a variety of other systems [75,77–80].

In general, the systematic expansion can be applied to any lattice, the crucial part is the choice of building block, which builds up the infinite lattice by translational symmetries t . All generated configurations are extended by adding the building block in each step of the expansion. There are two main aspects that need to be considered for the choice of building block. First, the number of generated clusters scales superexponentially with the order of the expansion. Second, the complexity of solving these clusters scales exponentially with systems size, which limits the maximal size of practically solvable clusters. Choosing a building block with a large number of sites induces a relatively small number of clusters which are still solvable; a building block with a small number of sites induces a very large number of solvable clusters to a degree that one may not be able to reach the largest solvable cluster size. Common choices of the building block in the square lattice is a single site or a complete square. Physically motivated, most NLCE approaches in the pyrochlore lattice use a tetrahedra expansion, based on clusters of complete tetrahedra, such that no dangling spins or triangles occur. In the following we discuss the cluster expansion for multisite unit cells and compare expansions in the pyrochlore lattice based on three different building blocks, in particular the *single site expansion*, the *unit cell expansion*, and the *tetrahedra expansion*. The latter turns out to yield the most reliable

results and represents the optimal expansion in that our results include full exact diagonalization of all clusters consisting of up to eight tetrahedra. The largest included clusters thus consist of 25 spins $\frac{1}{2}$, which host crucial loops of six and eight spins in the lattice.

1. Basic recipe

NLCE generates all possible subclusters c (subject to the choice of building blocks) which are embedded in the infinite lattice structure \mathcal{L} and contribute to the thermodynamic limit $P(\mathcal{L})/N$ per site. The contribution of each cluster c is given by its weight $W_P(c)$, describing the new (i.e., not included at lower order) contribution of c to P , and its multiplicity $L(c)$, describing the number of possible embeddings of c in \mathcal{L} .

The generality of this idea allows the definition of various building blocks. For the pyrochlore lattice studied here we include (a) all clusters built from single lattice sites, (b) all clusters built from complete (tetrahedral) unit cells, and (c) all clusters built from complete tetrahedra.

All possible configurations of clusters at order n , subject to the choice of building blocks, are expanded further in each step of NLCE from $n \rightarrow n + 1$ by adding building blocks in all possible ways. The initial cluster given by the building block (or multiple clusters given by inequivalent building blocks) have to respect translational symmetries t and cover the whole infinite structure \mathcal{L} by applying these symmetries t . It is important to point out the crucial difference between translational t and nontranslational symmetries s such as rotation or reflection. Translational symmetries are used to compare possible configurations of clusters. The systematic NLCE expansion considers only clusters that are independent/inequivalent by translational symmetries since the number of embeddings in \mathcal{L} is counted *per lattice site*. Two clusters are equivalent if translational symmetries map one cluster to the other. Hence, translational symmetries are used to compare all generated configurations of clusters and ensures that each cluster is only used once in the expansion. On the other hand, nontranslational symmetries can be applied to obtain a reduction in the number of clusters, which is however not mandatory for the correct result.

Each step of the expansion (at expansion order n) generates a set of connected clusters C_n . This large set can be reduced to the set of clusters which are not related by lattice symmetries S_n , and subsequently to topologically distinct clusters T_n of size n (number of building blocks). The set of connected clusters C_n includes all possible clusters which are embedded in \mathcal{L} that are not related by any translational symmetry t to each other. The set of clusters not related by lattice symmetries S_n are given by a subset of C_n ; each cluster in S_n is neither related by translational t nor nontranslational symmetry s to each other. Applying all nontranslational symmetries s to S_n generates all connected clusters C_n . Moreover, the set S_n can be reduced further. Even though clusters are not related by any symmetry, they can exhibit the same interaction topology and hence, generate the same Hamiltonian matrix. We describe each cluster by its interaction graph G , where its nodes $i \in N_G$ correspond to the spins included in the cluster and its edges $(i, j) \in E_G$ correspond to nearest-neighbor interaction terms of the Hamiltonian (1). Two clusters are

topologically equivalent if there is a graph isomorphism $\pi : N_{G_1} \rightarrow N_{G_2}$ (bijective) mapping G_1 on G_2 while preserving its structure; that means if $(i_1, j_1) \in E_{G_1}$ is an edge of G_1 , then $[\pi(i_1), \pi(j_1)] \in E_{G_2}$ needs to be an edge of G_2 . Hence, the set of topologically distinct clusters \mathcal{T}_n is a subset of \mathcal{S}_n . Using building blocks including more than one lattice site (like the tetrahedron) requires us to check the topological structure on the full connectivity graph including all sites.

The multiplicity $L(c)$ assigned to each cluster c describes the number of possible embeddings in the infinite structure \mathcal{L} . All possible clusters (subject to the choice of building blocks) of size n are included in \mathcal{C}_n ; hence, the multiplicity of each cluster is one. Nontranslational symmetries s reduce the set of connected cluster to \mathcal{S}_n and summarize multiple clusters in \mathcal{C}_n to one representative cluster $c \in \mathcal{S}_n$. The multiplicity of each cluster in \mathcal{S}_n is given by the number of nontranslational symmetries s that transform the cluster to another cluster that is not related to the first one by any translation t . Again, multiple clusters in \mathcal{S}_n are summarized to one representative cluster $c \in \mathcal{T}_n$, its multiplicity given by the number of topologically equivalent clusters. Hence, the multiplicity of topologically invariant clusters is simply the sum of the multiplicities of all topologically equivalent clusters in \mathcal{S}_n . Summing all multiplicities of clusters in \mathcal{S}_n or \mathcal{T}_n equals the number of connected clusters:

$$\sum_{c \in \mathcal{S}_n} L_{\text{sym}}(c) = \sum_{c \in \mathcal{T}_n} L_{\text{top}}(c) = |\mathcal{C}_n|. \quad (14)$$

For clarity we will drop the index top in what follows, i.e., $L(c) = L_{\text{top}}$, etc. The basic recipe to expand the clusters by one building block ($n \rightarrow n + 1$) is equivalent for all geometries and building blocks:

(i) Starting from all clusters not related by lattice symmetries of size n in \mathcal{S}_n , we generate new clusters by adding a building block to every free nearest neighbor. Again, we only consider clusters that are distinguishable by translational symmetries t .

(ii) Nontranslational symmetries s are used to reduce the set of all newly generated expansions to create the set \mathcal{S}_{n+1} . Applying all nontranslational symmetries s to \mathcal{S}_{n+1} generates the full set of connected clusters \mathcal{C}_{n+1} with respect to translational equivalence.

(iii) Clusters in \mathcal{S}_{n+1} are further reduced to obtain all topologically distinct clusters in \mathcal{T}_{n+1} .

Each cluster in \mathcal{T}_n contributes to the expansion according to its multiplicity $L(c)$ and weight $W_P(c)$. The n th order of NLCE is given by

$$P(\mathcal{L})/N|_n = \sum_{m=0}^n \sum_{c \in \mathcal{T}_m} L(c)W_P(c). \quad (15)$$

The weight assigned to a cluster c is defined with respect to all smaller subclusters $s \subset c$ (subject to the choice of building blocks) which can be embedded in c ; hence, it extracts contributions of c to P which are not covered by smaller clusters:

$$W_P(c) := P(c) - \sum_{s \subset c} W_P(s). \quad (16)$$

In practice, the thermodynamic observable P needs therefore to be calculated for all topologically invariant clusters. The

TABLE I. Nontranslational symmetries in the pyrochlore lattice.

s_1 : id	s_7 : $2(0, \bar{y}, y)$
s_2 : $3^+(x, x, x)$	s_8 : $2(\bar{x}, 0, x)$
s_3 : $3^-(x, x, x)$	s_9 : $2(x, \bar{x}, 0)$
s_4 : $m(x, y, y)$	s_{10} : $\bar{1}(0, 0, 0)$
s_5 : $m(x, y, x)$	s_{11} : $\bar{3}^+(x, x, x; 0, 0, 0)$
s_6 : $m(x, x, z)$	s_{12} : $\bar{3}^-(x, x, x; 0, 0, 0)$

extensive property of P induces a zero weight for disconnected clusters [32] which do not have to be considered in (16). Expanding clusters only by nearest neighbors guarantees connected clusters in \mathcal{C}_n , \mathcal{S}_n , and \mathcal{T}_n . The recursive definition of the weight ensures the convergence towards the infinite structure \mathcal{L} which is the thermodynamic limit.

It is not necessary to use all (or any) nontranslational symmetries in the NLCE. A lower number of nontranslational symmetries increases the number of clusters $|\mathcal{S}_n|$ such that each cluster has a lower multiplicity. In fact, it is possible to consider the identity id as the only nontranslational symmetry, then $\mathcal{C}_n = \mathcal{S}_n$. However, checking the topological structure of these clusters generates the same set of topologically distinct clusters \mathcal{T}_n . The computational effort can increase drastically with a low number of symmetries, since the number of clusters grows exponentially.

2. Pyrochlore lattice and building blocks

The underlying Bravais lattice is given by a fcc structure with a tetrahedral unit cell (four sites). Hence, in order to respect the translations, all expansions we use focus on the tetrahedral unit cell and converge to the thermodynamic limit per unit cell and not per site, thus accounting for an extra factor of four.

Its symmetries are described by the space group $Fd\bar{3}m$ (227); it contains 192 symmetry operations. However, only 12 symmetries are purely nontranslational and used in the NLCE, see Table I.

s_2 , s_3 and s_7 , s_8 , s_9 describe threefolded and twofolded rotations, respectively. Reflections are given by s_4 , s_5 , s_6 , and s_{10} represents the inversion. s_{11} and s_{12} combine a threefold rotation with the inversion.

As discussed earlier, the choice of building block is crucial. In principle, various geometries (such as dimers, hexagons, or multiple tetrahedra) can be used as building blocks as long as they respect the translations t and cover the whole lattice. Each expansion we use is embedded in the fcc structure of the pyrochlore with either equivalent (unit cell expansion) or multiple inequivalent (single site and tetrahedra expansion) building blocks.

An efficient implementation of translational symmetries is essential due to the exponentially increasing complexity. Labeling lattice sites along the translational axes automatically describes the translational symmetries t by a simple index shift.

a. Single site expansion: As pointed out earlier, the single site expansion generates a large number of clusters of relatively small size. The advantage of this approach is its complete generality. In contrast to the single site expansion in Bravais lattices (e.g., square or triangular lattice [31,32,74]),

the unit cell in the pyrochlore lattice consists of four sites, which need to be treated inequivalently. Hence, the starting point of the single site expansion are four sites arranged in the unit cell/tetrahedron, which covers the whole lattice by translations. Applying translation symmetries to these four sites generates the full pyrochlore lattice. Here all symmetries in Table I can be applied to find clusters which are related by lattice symmetries to reduce the complexity.

b. Unit cell expansion: The unit cell expansion is related to the single site expansion in the fcc lattice, substituting each site in the obtained clusters by the tetrahedral unit cell. However, working within the pyrochlore lattice, the symmetries in Table I have to be examined more carefully: We require the symmetries to preserve the unit cell structure such that only entire unit cells are mapped to each other. Only the symmetries s_1 to s_6 in Table I, which are a subset of the fcc-lattice symmetries, are unit cell conformal. As mentioned before, working with a lower number of symmetries produces the same results. Since the building block includes more than one site, the topological structure has to be compared on the level of the full connectivity graph including all lattice sites. The advantage of this approach is the consideration of larger clusters due to a much slower growth of the number of clusters with the number of unit cells. However, the connection between the unit cells are dangling bond that do not reflect the geometrical frustration. In the presence of magnetic fields, the Hamiltonian has bond and site terms. Therefore, we include a single site (yielding the simplest contribution of the site terms) as the 0th order in the expansion ($m = 0$) in (15), which is embedded four times in the unit cell.

c. Tetrahedra expansion: The central motif of the pyrochlore lattice is the tetrahedron. An examination of the lattice shows that there are two types of tetrahedra: up pointing tetrahedra (these are the unit cells) and down pointing tetrahedra, which correspond to the interaction of each spin in the unit cell to three neighboring spins in *different* unit cells. Both the single site and unit cell expansion do not respect this structure, leading to dangling bonds in the case of most clusters in the single site expansions and incomplete down pointing tetrahedra in the case of the unit cell expansion.

For this reason we use an expansion including all clusters with *complete tetrahedra* [33–37,76]. The systematic expansion focuses on an hourglass structure composed by two inequivalent building blocks of tetrahedra (up/down pointing) which are placed in the underlying fcc lattice and expands these as described before. Our comparison of results for the heat capacity and the magnetic susceptibility demonstrates that this intuition is correct and that this expansion is indeed superior (cf. comparison in Appendix A).

It is important to note that the physical size (number of spins) of each cluster is not uniquely related to the order of the expansion, due to an overlap of up- and down-facing tetrahedra. This means that at an expansion order (given by the total number of tetrahedra) n_{tetra} clusters of different sizes are included. Similarly to the unit cell expansion, this expansion leads to a relatively small number of large clusters. Again, we need to consider the 0th order contribution of a single site. We did not apply any symmetry due to the small number of edges of each tetrahedron and the low order of expansion and rely on the automatic identification of topologically

equivalent clusters by directly comparing their interaction graphs.

In Appendix B we provide a comparison of the number of clusters generated at each order in the three expansions discussed here. We list the number of connected $|C_n|$, not related by lattice symmetries $|\mathcal{S}_n|$ (if present), and topologically distinct clusters $|\mathcal{T}_n|$ in Tables II, III, and IV. The visualization of these results is shown in Fig. 12, clearly showing that the number of clusters with at most n sites is smallest in the tetrahedra expansion, leading to a tractable number at the edge of feasibility of full diagonalization (the maximal symmetry block dimension of the Hamiltonian is 228592 for clusters with 25 sites (full Hilbert space 3.355×10^7)). Additionally, all topologically distinct clusters of the tetrahedra expansion can be found in Appendix B.

3. Resummation algorithms

Correlations are increasingly long ranged as temperature is lowered. Hence, contributions of larger cluster have to be taken into account to converge to the thermodynamic limit. Accessing these orders is limited due to the exponentially increasing complexity regarding the number of clusters or Hilbert space dimension.

One effective tool to obtain reliable data for lower temperature are resummation algorithms like Euler's transformation [81] which can accelerate the convergence of NLCE. Detailed descriptions and examples can be found in Refs. [30,32,74,75,77–80].

Resummation algorithms rely on a systematic usage of lower orders of the series and are most effective if many terms are included. They are guaranteed to converge to the limiting value of the underlying series, or not at all. In this work we use the Euler transform of our NLCE data for the expansion up to $n = 8$ tetrahedra and compare the highest order Euler transform of the Euler transform up to expansion orders $n = 7$ (and also $n = 6$), to ensure convergence of our results.

Euler's transformation is particularly useful for alternating series, which are transformed according to [81]

$$\sum_{s=0}^{\infty} (-1)^s u_s = \sum_{s=0}^{\infty} \frac{(-1)^s}{2^{s+1}} [\Delta^s u_0], \quad (17)$$

where Δ^s is the s fold application of the forward difference operator Δ , defined by $\Delta u_n := u_{n+1} - u_n$. Euler's method can be derived by repeated application of summation by parts [82].

In the present work we find that the Euler transformation of our NLCE series up to $n = 8$ tetrahedra (i.e., containing eight terms) yields a *significant improvement of convergence* at low temperatures. We always compare the Euler transformation of the first $n = 7$ and $n = 8$ terms of the series to ensure that the results are indeed converged, yielding reliable results down to $T \approx 0.25$ for thermodynamic properties of the $S = \frac{1}{2}$ pyrochlore antiferromagnet as shown in detail in what follows.

D. Exact diagonalization of clusters

1. Cluster symmetries from graph automorphisms

The numerical linked cluster expansion (NLCE) expresses thermodynamic observables as a series expansion in terms of the exact solution of a large number of finite size clusters.

Since the number of clusters grows factorially with the number of constituents, it is useful to use an automatic strategy to identify cluster symmetries, which are used for block diagonalizing the Hamiltonian. Here we provide a practical description of the method with only minimal reference to graph and group theory to make it accessible. A similar pedagogical description for the exploitation of translational symmetry can be found in Ref. [83]. In all of the following discussion we use the computational S^z basis, in which each basis state $|\sigma_1, \sigma_2, \dots\rangle$ is an eigenstate of all local S_i^z operators and labeled by their eigenvalues σ_i .

As described in the NLCE part, we identify a finite cluster with its interaction graph G , where its nodes $i \in N_G$ correspond to the spins and its edges $(i, j) \in E_G$ correspond to nearest-neighbor interaction terms of the Hamiltonian (1). Hence, the Hamiltonian is defined by

$$H = \sum_{(i,j) \in E_G} \vec{S}_i \cdot \vec{S}_j + h \sum_{i \in N_G} S_i^z. \quad (18)$$

The sum $(i, j) \in E_G$ runs over all edges (i, j) of the graph G , the sum $i \in N_G$ runs over all sites. We notice that any *automorphism* of the graph G leaves the Hamiltonian invariant, since an automorphism is a *permutation* π of nodes which maps the graph onto itself, such that for any edge $(i, j) \in E_G$, the mapped edge has to be in G as well: $[\pi(i), \pi(j)] \in E_G$.

This means that graph automorphisms are *symmetries* of the Hamiltonian and commute with it, which implies that we can simultaneously diagonalize the automorphism and the Hamiltonian.

The matrix representation \mathbf{A} of a graph automorphism A (which necessarily is a permutation π_A of graph nodes) can be obtained by noticing that any basis state is transformed as

$$A|\sigma_1, \sigma_2, \dots, \sigma_N\rangle = |\sigma_{\pi_A(1)}, \sigma_{\pi_A(2)}, \dots, \sigma_{\pi_A(N)}\rangle. \quad (19)$$

\mathbf{A} is a permutation matrix on the set of basis states and $\mathbf{A}^T \mathbf{A} = \mathbf{1}$, i.e., \mathbf{A} is orthogonal with eigenvalues on the complex unit circle. Since \mathbf{A} is a permutation matrix, it is *idempotent with a certain order* $N_A < |N_G|$: $\mathbf{A}^{N_A} = \mathbf{1}$. Therefore, its eigenvalues are given by the N_A th roots of unity: $e^{i2\pi n_A/N_A}$, with $n_A \in \{0, \dots, N_A - 1\}$.

This means we can block diagonalize H into n_A blocks, labeled by the eigenvalues $e^{i2\pi n_A/N_A}$ of \mathbf{A} . We denote the order $o(\pi)$ of an automorphism π by the minimal number $n \in \mathbb{N}$ such that $\pi^n = \mathbf{1}$.

2. Identification of largest commuting automorphism subgroup

Typical interaction graphs G have a large number of (independent) graph automorphisms. However, typically *not all of them commute!* Since only commuting graph automorphisms can be used together to further reduce the block size of the Hamiltonian, we want to find the largest Abelian (commuting) subgroup of the complete automorphism group of G .

In order to do so, we start by generating the automorphism group of the graph. In the next step we check for each pair of automorphisms if they commute. This can be interpreted by a new graph C , in which each automorphism of G corresponds to a node and two nodes are connected if and only if the corresponding automorphisms commute. What we are looking for is a subgroup $\mathcal{U} \subset N_C$ in which each node in

\mathcal{U} is connected to all other nodes of the subgroup. Hence, each element in \mathcal{U} commutes with each other. This is called a *clique* in graph theory. Finding the largest Abelian subgroup of the automorphism group is therefore identical to finding the largest clique in C . In general, the largest clique is not uniquely determined.

After determining the largest clique \mathcal{U} of automorphisms, we need to identify a *minimal* set \mathcal{H} of independent generators of the Abelian subgroup. Each element of the clique \mathcal{U} has to be generated *uniquely* by elements of \mathcal{H} . Assume \mathcal{H} includes $m := |\mathcal{H}|$ elements $h \in \mathcal{H}$ of order $o(h)$, then there is exactly one tuple $(n_{h_1}, \dots, n_{h_m})$ of integers for each element u of the subgroup, which generates u from the corresponding integer powers of the generators:

$$\forall u \in \mathcal{U} : \exists! (n_{h_1}, \dots, n_{h_m}) \in \mathbb{N}^m \quad \text{with} \quad 0 \leq n_{h_i} < o(h_i),$$

$$\text{such that} \quad u = \prod_{h \in \mathcal{H}} h^{n_h}. \quad (20)$$

The product in EQ. (20) refers to the composition of permutations. The ordering is arbitrary since all generators commute. Each element in \mathcal{U} is represented uniquely by a multi-index $(n_{h_1}, \dots, n_{h_m}) \in \mathbb{N}^m$ where each number is smaller than the corresponding order of the generator. This multi-index determines a phase in the symmetrized basis (cf. Sec. III. D. 3). The bijective mapping from \mathcal{U} to \mathbb{N}^m , respecting the order $o(h_i)$ of the generators h_i , induces the following relation between the cardinality of the Abelian subgroup $|\mathcal{U}|$ and the orders of its generators:

$$|\mathcal{U}| = \prod_{h \in \mathcal{H}} o(h). \quad (21)$$

In practice we create the minimal set of generators by starting with the elements exhibiting the highest order in the commuting subgroup \mathcal{U} . First, one element with the highest order will be added in \mathcal{H} . A new element g is added to \mathcal{H} if it does not violate the bijective mapping described through (20). That is, all possible compositions $g^n \circ h$, for $n = 0, \dots, o(g) - 1$ and $h \in \mathcal{H}$, generate uniquely a new element in the subgroup \mathcal{U} . The generating set \mathcal{H} is not uniquely determined.

In Fig. 2 we show an illustration for an example interaction graph (hourglass composed of two corner sharing tetrahedra, corresponding to the $n = 2$ cluster in NLCE) along with its 71 nontrivial automorphisms and their commutation relations. The largest clique for this example has eight nontrivial automorphisms with two generators A and B [e.g., given by the node permutations $\pi_A = (2, 0, 1, 3, 5, 6, 4)$, $\pi_B = (2, 0, 1, 3, 6, 4, 5)$], generating independent C_3 rotations of the upper and lower tetrahedron. In the present work we rely on established algorithms for the identification of graph iso- and automorphisms bundled in the package *nauty* [84] and used the clique maximization algorithm described in Ref. [85].

3. Symmetrized basis

Once we have obtained a set of independent generators \mathcal{H} and the multi-indices referring to the largest Abelian subgroup \mathcal{U} of the graph automorphism group, we can proceed with the block diagonalization of the Hamiltonian. A subgroup of size $|\mathcal{U}|$ induces the same number of blocks, each is uniquely iden-




FIG. 2. Left: Example interaction graph G of an hourglass composed by two corner sharing tetrahedra. Right: Resulting automorphism graph C , where each node represents one automorphism, and commuting automorphisms are connected by an edge. The red colored nodes represent the largest clique in the graph, in which all automorphisms commute (red edges). The identity $(0,1,2,3,4,5,6)$ commutes with all automorphisms and has therefore been omitted from the graph C . It is part of the largest clique. The largest clique consists of the following eight nontrivial permutations: $(2, 0, 1, 3, 5, 6, 4)$, $(1, 2, 0, 3, 6, 4, 5)$, $(2, 0, 1, 3, 6, 4, 5)$, $(0, 1, 2, 3, 6, 4, 5)$, $(1, 2, 0, 3, 4, 5, 6)$, $(2, 0, 1, 3, 4, 5, 6)$, $(1, 2, 0, 3, 5, 6, 4)$, $(0, 1, 2, 3, 5, 6, 4)$, which are marked in red, the blue edges indicate a smaller clique with five nontrivial permutations.

tified by m (number of minimal generators) quantum numbers given by multi-indices described in (20). Each index refers to the phase of the eigenvalue of the corresponding generator; the commutation relations of the generators allows the simultaneous diagonalization of all generators and the Hamiltonian. Each computational basis state $|\vec{\sigma}\rangle = |\sigma_1, \dots, \sigma_N\rangle$ has to be replaced by a symmetrized state induced by the quantum numbers $\varepsilon = (\varepsilon_1, \dots, \varepsilon_m) \in \mathbb{N}^m$ which is given by

$$|\sigma_1, \dots, \sigma_N; \varepsilon\rangle = \frac{1}{\sqrt{\mathcal{N}_{\vec{\sigma}}^{\varepsilon}}} \sum_{u \in \mathcal{U}} \prod_{i=0}^m e^{i2\pi n_i^u \varepsilon_i / o(h_i)} u |\sigma_1, \dots, \sigma_N\rangle, \quad (22)$$

where $\mathcal{N}_{\vec{\sigma}}^{\varepsilon}$ is the normalization constant of the state, $n^u = (n_1^u, \dots, n_m^u) \in \mathbb{N}^m$ is the multi-index referring to $u \in \mathcal{U}$ defined by (20), and $o(h_i)$ is the order of the generator $h_i \in \mathcal{H}$. The eigenvalue of each generator $h_i \in \mathcal{H}$ is obtained (using the properties of the generators and the complex roots of unity) from

$$h_i |\sigma_1, \dots, \sigma_N; \varepsilon\rangle = e^{i2\pi \varepsilon_i / o(h_i)} |\sigma_1, \dots, \sigma_N; \varepsilon\rangle. \quad (23)$$

It is important to note that multiple unsymmetric basis states $|\vec{\sigma}\rangle$ typically generate the *same* (apart from a phase) symmetric state. This is in fact true for any basis state which is in the set

$$F_{\vec{\sigma}} = \{u|\vec{\sigma}\rangle \mid u \in \mathcal{U}\}, \quad (24)$$

which we call the ‘‘family of the state $|\vec{\sigma}\rangle$ ’’ generated by the commuting subgroup \mathcal{U} of the graph automorphism group. Therefore, each symmetric basis state has to be added *only once* to the symmetric basis. This is typically ensured by using a *parent state* of the family, for example the state $|\vec{\sigma}\rangle$ with the

lowest binary representation; this parent state is denoted by $p(F_{\vec{\sigma}})$.

Crucially, some unsymmetric basis states $|\vec{\sigma}\rangle$ *do not generate any symmetric state* in a given symmetry sector. This happens if the basis state is *incompatible with the symmetry sector* and the sum of phase factors cancels, leading an unnormalizable state. An example is the state $|\uparrow\uparrow\uparrow \dots \uparrow\rangle$. For any graph and any automorphism u , it is mapped to itself. Therefore, its family is $F_{\uparrow\uparrow\uparrow \dots \uparrow} = \{|\uparrow\uparrow\uparrow \dots \uparrow\rangle\}$. In the symmetry sector $0^m := (0, \dots, 0)$, we obtain $|\uparrow\uparrow \dots \uparrow; 0^m\rangle = |\uparrow\uparrow \dots \uparrow\rangle$. In all other sectors, however, we get $|\uparrow\uparrow \dots \uparrow; n \neq 0^m\rangle = 0$, i.e., this state does not appear in other sectors. As in the general example, this mechanism leads to an imbalance of the size of sectors, the sector 0^m always being the largest.

It is crucial to ensure the correct normalization of the symmetric basis states $|\vec{\sigma}\rangle$ and $|\vec{\sigma}'\rangle$ introduced in Eq. (22). We require

$$\langle \vec{\sigma}; n \mid \vec{\sigma}'; n' \rangle = \delta_{nn'} \delta_{p(F_{\vec{\sigma}}), p(F_{\vec{\sigma}'})}, \quad (25)$$

i.e., states are orthogonal if they correspond to different parents (and therefore families), or if they correspond to different symmetry sectors n .

We note that in addition to the graph automorphisms, the Heisenberg model we study has additional spin symmetries. Here we exploit the conservation of the total z component of the spin, because $[\sum_i S_i^z, H] = 0$. Since all computational basis states we use are already eigenstates of the total z component $\sum_i S_i^z$, and since $\sum_i S_i^z$ also commutes with the graph automorphisms, this symmetry is trivial to exploit: A simple reordering of basis states by their z magnetization makes the Hamiltonian block diagonal. Additionally, we exploit the spin inversion symmetry $[Q, H] = 0$ with $Q := \prod_i S_i^x$ in the sector $m_z = 0$. The spin inversion is also used to deduce the results for the $-S_z$ sector from an already solved S_z sector.

4. Hamiltonian submatrix in symmetry sectors

Before we can fully diagonalize the Hamiltonian, we need to represent each block of H (labeled by the quantum numbers ε) in the symmetric basis. For simplicity we only focus on the graph automorphisms and ignore symmetries defined by S_z and Q (which are a trivial extension). Hence, we need to construct the matrix elements $\langle \vec{\sigma}; \varepsilon \mid H \mid \vec{\sigma}'; \varepsilon \rangle$; note that by construction the interblock matrix elements $\varepsilon \neq \varepsilon'$ are zero.

Let us apply the Hamiltonian to a symmetrized basis state, exploiting the fact that H commutes with $u \in \mathcal{U}$:

$$H |\vec{\sigma}; \varepsilon\rangle = \frac{1}{\sqrt{\mathcal{N}_{\vec{\sigma}}^{\varepsilon}}} \sum_{u \in \mathcal{U}} \prod_{i=0}^m e^{i2\pi n_i^u \varepsilon_i / o(h_i)} u H |\vec{\sigma}\rangle. \quad (26)$$

The Hamiltonian is expressed as a sum of nonbranching terms: $H = \sum_b h_b$; note that the permutations $u \in \mathcal{U}$ commute with each single term h_b . The operators h_b can be divided into diagonal operators (which do not change the state $|\vec{\sigma}\rangle$) and off-diagonal operators, which yield a different state $|\vec{\sigma}_b\rangle$ together with the corresponding matrix element by $h_b |\vec{\sigma}\rangle = h_b^{\vec{\sigma}\vec{\sigma}_b} |\vec{\sigma}_b\rangle$. $|\vec{\sigma}_b\rangle$ is in general not a parent state. It is however related to its parent state by a symmetry operation $|\vec{p}_b\rangle = u_0 |\vec{\sigma}_b\rangle = p(F_{\vec{\sigma}_b})$. Note that u_0 is not determined uniquely; multiple permutations can fulfill this mapping. Also, the new state is not necessarily a valid state in the symmetry sector ε ,

in which case it will be canceled by later terms. The parent state is assigned to an index of the basis; the referring matrix element can be calculated as follows:

$$\begin{aligned}
& \langle \bar{p}_b; \varepsilon | h_b | \bar{\sigma}; \varepsilon \rangle \\
&= \langle \bar{p}_b; \varepsilon | \frac{1}{\sqrt{\mathcal{N}_{\bar{\sigma}}^\varepsilon}} \sum_{u \in \mathcal{U}} \prod_{i=0}^m e^{i2\pi n_i^\varepsilon \varepsilon_i / o(h_i)} u h_b | \bar{\sigma} \rangle \\
&= \frac{1}{\sqrt{\mathcal{N}_{\bar{\sigma}}^\varepsilon \mathcal{N}_{\bar{p}_b}^\varepsilon}} \sum_{u, u' \in \mathcal{U}} \prod_{i=0}^m e^{i2\pi (n_i^\varepsilon - n_i^{\varepsilon'}) \varepsilon_i / o(h_i)} \langle \bar{p}_b | u' u | \bar{\sigma}_b \rangle h_b^{\bar{\sigma} \bar{\sigma}_b}.
\end{aligned} \tag{27}$$

Terms in Eq. (27) only have a nonzero contribution if and only if $|\bar{\sigma}_b\rangle$ is mapped by $u'u$ to its parent state $|\bar{p}_b\rangle$. The matrix elements of the Hamiltonian in the symmetrized basis are therefore given by the unsymmetrized matrix elements $h_b^{\bar{\sigma} \bar{\sigma}_b}$ multiplied by symmetry sector dependent phase factors and normalization constants. We note in passing that the sums over phase factors can yield the normalization constant $\mathcal{N}_{\bar{p}_b}^\varepsilon$ [83].

IV. RESULTS

Using the combination of the methods described in Sec. III, we address thermodynamic properties of the pyrochlore quantum Heisenberg antiferromagnet. We start by considering the SU(2) symmetric case without an applied magnetic field and present our results for the heat capacity (Sec. IV A), the magnetic susceptibility (Sec. IV B), and the thermodynamic entropy. We compare the results for different orders in the numerical linked cluster expansion (NLCE) and their Euler transform to show that the high temperature regime down to $T \approx 0.25$ is converged to the thermodynamic limit. We furthermore compare the NLCE results to the solution of finite size clusters obtained using canonical typicality and finite temperature DMRG.

We also present results obtained from DMRG for the static spin structure factor at finite temperature (Sec. IV C), as well as NLCE results for the heat capacity and the magnetization at finite applied magnetic field h (Sec. IV D).

Figure 3 shows by example of the specific heat how the converged results were obtained. The blue curve shows the heat capacity for a finite size cluster with $N = 32$ sites (8 unit cells, inset) in comparison with the results of the numerical linked cluster expansion for different orders n (top panel), indicating the number of complete tetrahedra in the clusters included in the expansion. To accelerate the convergence of the NLCE series at lower temperatures, we apply the Euler transformation up to order n in the series (cf. Sec. III. C. 3). We have furthermore calculated the specific heat capacity for a larger cluster with $N = 48$ sites and periodic boundary conditions using our SU(2) symmetric finite temperature DMRG method with finite bond dimensions ranging from $\chi = 2000$ to $\chi = 10000$. These results were obtained from a numerical derivative of the (spline interpolated) energy as a function of inverse temperature. For $T > 2$, the results are converged with bond dimension and agree with the NLCE and finite $N = 32$ cluster results (bottom panel). At lower temperatures, the dependence of the results on the bond dimension becomes significant and we extrapolate to infinite bond dimensions




FIG. 3. Comparison and convergence of heat capacity of the spin- $\frac{1}{2}$ pyrochlore Heisenberg model determined via different methods. Top: Numerical linked cluster expansion (NLCE) with clusters of $n = 2, \dots, 8$ complete tetrahedra, as well as their Euler transform for NLCE orders $n = 6, 7$, and 8 (red curves). These results appear to be converged down to temperatures of about $T = 0.25$. The blue curve shows the result for a 32 site (8 unit cells) cluster with periodic boundaries obtained from canonical typicality. The error bars (shaded blue area) reflect variations in results obtained from sampling over different random vectors. Bottom: DMRG results obtained from the numerical (spline) derivative of the energy $(H)_\beta$ for different bond dimensions χ obtained by finite temperature DMRG (purification) and the blue crosses represent their extrapolation to $\chi \rightarrow \infty$ using a quadratic polynomial in $1/\chi$. The error bars indicate the distance of the extrapolated from the $\chi = 10000$ results. Typicality and Euler transform of NLCE results as in top panel.

using a quadratic polynomial in $1/\chi$ yielding a very good match with the $N = 32$ and NLCE results within the accessible temperature range ($T > 0.62$). The error bars indicate the distance of the extrapolated value from the largest bond dimension ($\chi = 10000$).

A. Heat capacity in zero field

The specific heat capacity quantifies the change of the internal energy as a function of temperature and is directly accessible in experiment. Our results and a comparison to experimental results are summarized in Fig. 1.

Starting at high temperatures $T \gg 1$, the heat capacity decays as $1/T^2$, as required in the leading order high temperature expansion. In the regime down to $T = 2$, all orders of the NLCE with $n > 5$ agree with each other, and also with the

finite size $N = 32$ result from typicality. The Euler transform of our NLCE results agrees between orders $n = 7$ and $n = 8$ down to $T \approx 0.25$, which we take to indicate that the series is converged over this range.

Crucially, this allows us to resolve unambiguously the maximum of the heat capacity located at $T \approx 0.57$. In the proximity of the maximum, the results for the finite size cluster $N = 32$ deviate significantly from the NLCE, which indicates that in the regime $T < 2$ the correlations beyond the size of the $N = 32$ cluster start playing a discernible role. From the strong dependence of the DMRG results on the bond dimension around the location of the maximum of the heat capacity, as well as from the visible discrepancy of the specific heat for finite size clusters compared to the converged NLCE results close to the maximum, we conclude that the system enters a nontrivial quantum regime at temperatures $T \approx 1$, where it exhibits entanglement beyond what is representable faithfully by $\chi = 10000$ matrix product states.

The heat capacity of the $N = 32$ cluster exhibits a second maximum at low temperatures, similarly to what was observed previously in a different pyrochlore model [86]. However due to the divergence from converged results of the NLCE in this regime, which is not subject to finite-size effects of this kind, we conclude that this feature is likely not representative for the thermodynamic limit.

Indeed, the converged part of our NLCE data shows a rapid decrease of the heat capacity as T is lowered from the maximum. Therefore, if there is an additional feature, it must be well separated from the maximum we have found.

In order to gain further insight into the low- T regime, we have calculated the thermodynamic entropy as a function of temperature,

$$S(T_2) - S(T_1) = \int_{T_1}^{T_2} dT \frac{C_V}{T}, \quad (28)$$

with $S(\infty) = \ln 2$. A direct calculation of the entropy per site in NLCE using Eq. (6) shown in Fig. 4 indeed agrees with this temperature integral down to the lowest temperatures for which our (Euler transformed) NLCE is converged.

Interestingly, we find that just over half the total entropy is released down to $T \approx 0.25$, where the entropy is $S/N \approx 0.33 \approx 0.47 \ln 2$. This in turn means that the spectral weight below the maximum is huge, and there is plenty of scope for further interesting behavior, the nature of which we are unfortunately unable to determine from our approach. (In several classical models [6], as well as some fine-tuned quantum models [87–89], not all the entropy is released even at $T = 0$, but in real systems, the third law of thermodynamics stipulates that this is not what actually happens. For instance, for the finite size $N = 32$ cluster, a steep decrease at low $T \gtrsim 0.1$ is associated with its low- T peak in the specific heat.)

We will return to more detailed comparisons of this behavior with other models and experimental systems in the final discussion, Sec. V.

B. Magnetic susceptibility at $h = 0$

We consider the magnetic susceptibility χ/N [90] per lattice site in Fig. 5 as a function of temperature. As in the case of the heat capacity, we perform NLCE calculations up




FIG. 4. NLCE results for the thermodynamic entropy per lattice site S/N as a function of temperature T . The yellow through dark blue curves show raw NLCE data for different expansion order n up to $n = 8$ tetrahedra. The brown, black, and red curves are the corresponding Euler transform, showing converged entropies down to $T \approx 0.25$ for $n = 8$. These results agree with the temperature integral (light blue curve) of C_V/T of the Euler $n = 8$ data for the specific heat from Fig. 3. We also show the entropy obtained from the finite size 32 site cluster obtained from canonical typicality in the black curve.

to clusters of eight tetrahedra and apply the Euler transform to these results. Figure 5 shows the raw NLCE results at order $n = 8$ to be converged down to temperatures of about $T \approx 0.8$. The Euler transform improves the convergence of the series significantly, again down to a temperature of $T \approx 0.25$.

At high temperatures $T \gtrsim 1$, the susceptibility obtained by the various methods agrees: typicality for $N = 32$ site cluster, the $N = 48$ cluster result obtained in finite temperature DMRG, extrapolated to infinite bond dimension using a




FIG. 5. Magnetic susceptibility χ/N as a function of temperature T . We show raw data from different NLCE orders (yellow, green) as well as their Euler transform (red). The blue curve shows the result in the finite size cluster with $N = 32$ spins and the black crosses with error bars indicate the diagrammatic Monte Carlo results from Ref. [91]. We also show DMRG results for the $N = 32$ site cluster (blue crosses) and for the $N = 48$ site cluster (green crosses) extrapolated to infinite bond dimension, where the error bars indicate the distance of the extrapolated value to the largest bond dimension $\chi = 10000$.

quadratic polynomial in inverse bond dimension. At $T \approx 0.6$ the finite size magnetic susceptibility exhibits a pronounced maximum and decays rapidly to zero at lower temperatures. The Euler transform for the largest NLCE order clearly reveals a decrease of the magnetic susceptibility after a *maximum* at $T = 0.54$ in the thermodynamic limit.

We note that the magnetic susceptibility for very large system sizes was previously calculated in diagrammatic quantum Monte Carlo simulations in Ref. [91], corresponding to the black crosses with error bars in Figs. 5. These results agree with our NLCE and finite size results at temperatures above the maximum of the susceptibility. At low temperatures, however, they suggest a steady increase or plateau of the susceptibility instead of the maximum that we find. It would be desirable to push both our and the diagrammatic Monte Carlo method to higher orders in order to see which of the two apparently irreconcilable behaviors is the correct one.

C. Static spin structure factor at zero field

The static spin structure factor quantifies the spin correlation patterns present at a given T :

$$S(\vec{Q}) = \frac{4}{3N} \sum_{ij} \langle \vec{S}_i \cdot \vec{S}_j \rangle_{\beta} \cos \left[\vec{Q} \cdot (\vec{R}_i - \vec{R}_j) \right]. \quad (29)$$

Here we use finite temperature DMRG on the two clusters with 32 sites and full cubic symmetry and 48 sites with reduced symmetry, to investigate the static spin structure factor at finite temperature. Figure 6 shows the result in the 32 site cluster (top rows) and for the 48 site cluster (bottom rows), in the the (h, h, l) plane (i.e., $Q_x = Q_y$, rows 1, 3) and the $(h, l, 0)$ plane ($Q_z = 0$, rows 2, 4) of momentum space, extended over several Brillouin zones for different temperatures (columns). Both figures show the clear emergence of a correlation structure already at high T . This becomes more pronounced as T is lowered, without acquiring much additional structure: certainly, as expected for a highly frustrated magnet, no sharp Bragg peaks appear, which would have been indicative of magnetic ordering.

Indeed, with decreasing temperature, the weight in the center of the Brillouin zone ($\vec{Q} = 0$) decreases and moves to the boundary of the extended Brillouin zone, the most visible location of increasing intensity being at $(0, 0, \pm 4\pi)$.

At this location, alongside $(\pm 2\pi, \pm 2\pi, \pm 2\pi)$, one finds the incipient pinch points, well known from other pyrochlore magnets [7,13,15,24,42,43], as well as settings with an emergent U(1) gauge field more generally.

As seen in Fig. 6 top row, incipient bowties emerge in the spin correlation. These are a characteristic feature of the pyrochlore lattice antiferromagnets, reflecting the tendency to the formation of momentless tetrahedra. In contrast to the classical magnets, where individual tetrahedra can simultaneously be momentless, we observe a finite width of pinch points (within the limitations of our finite resolution), indicating the persistence of a nonzero net moment on the tetrahedra, reflecting the inability of the quantum system to host simultaneous singlet configurations of tetrahedra sharing a common spin.

Due to the low resolution in k space our finite system sizes up to $N = 48$ do not permit us to investigate more closely




FIG. 6. Static spin structure factor $S(\vec{Q})$ for different temperatures calculated with finite- T DMRG in the $N = 32$, upper rows ($N = 48$, lower rows) site cluster using bond dimension of $\chi = 8000$ ($\chi = 10000$). Even rows show the (h, h, l) plane in momentum space, odd rows the $(h, l, 0)$ plane. The DMRG ansatz for the two different clusters has a different symmetry.

how sharp the pinch points become, since the associated lengths at low temperatures may become longer than our linear cluster sizes. However, a *very* rough idea can be gleaned by considering the T dependence of the energy. The reason for this is that the energy simply encodes the total spin of each tetrahedron,

$$E = \sum_{\text{tet}} \frac{1}{2} \left(\sum_{i \in \text{tet}} \vec{S}_i \right)^2 + \text{const.} \quad (30)$$

The pinch points become infinitely sharp in the limit of vanishing tetrahedral magnetic moment $\sum_{i \in \text{tet}} \vec{S}_i = 0$ [42].

This condition is known to be met in classical Ising (where it is known as the ice rule [4–6]), XY, and Heisenberg models, but it cannot hold in the quantum realm, since a spin which is part of two tetrahedra cannot enter singlet bonds in both of them simultaneously. The deviation of the energy from the minimal value is hence a proxy for the possible sharpness of the pinch points. The total spin of a tetrahedron is plotted in Fig. 7 from our various methods. Eyeballing an extrapolation of this curve to $T = 0$, one finds $\sum_{i \in \text{tet}} \vec{S}_i \approx 1$, which confirms




FIG. 7. Total spin $\langle \sum_{\text{tet}} \vec{S}_i \rangle_\beta$ of a tetrahedron in the 32 (blue) site cluster as a function of temperature T . The red curves show the Euler transform of NLCE results and the black curve is the total spin of a single tetrahedron (four site cluster) for comparison.

that frustration precludes that all tetrahedra are in singlet states and hence a finite width of the pinch points; for an estimate of the pinch-point width based on PF-FRG, see Ref. [24].

D. Heat capacity and entropy in a magnetic field

We complement the above discussion by a further analysis of the behavior of the pyrochlore Heisenberg magnet in the presence of a finite field. For readability of our figures, we only show the converged part of eighth order NLCE results using the tetrahedra expansion. As before, we use the agreement of eighth and seventh order Euler transform as convergence criterion.

Figure 8 shows the heat capacity and entropy per site as a function of temperature for a range of different fields h applied in the [001] direction. We observe a shift of the maximum of the specific heat to higher temperatures at strong fields, as well as a nonmonotonic change of the height of the maximum.

An overall upward shift of the weight is not particularly surprising given the presence of an additional term in the Hamiltonian. Indeed, at high fields, the curve resembles an unspectacular paramagnet.

However, the structure of the low-energy spectral weight and its rearrangement at intermediate fields is complex. The weight at lowest energies is in large part due to nonmagnetic states which are not favored by the magnetic field. At the same time, the more numerous states with nonzero magnetization spread out as the field is applied. The failure of the Euler transform to converge to similarly low temperatures at intermediate fields is presumably due to a more complex behavior of the specific heat in this intermediate field regime.

E. Magnetization at $h > 0$

We finally consider the effect of a magnetic field on the magnetization per site m_z/N . At zero field there is no net magnetization in the absence of spontaneous symmetry breaking. Figure 9 shows converged NLCE results (with the highest order Euler transform) for the finite temperature magnetization (solid lines) in comparison with the result for a single tetrahedron (dashed lines). At high temperatures, these




FIG. 8. Magnetothermodynamics of the $S = \frac{1}{2}$ pyrochlore Heisenberg magnet. Top: heat capacity C_V/N per lattice site at varying fields as a function of temperature. We show data obtained from the Euler transform of the $n = 8$ tetrahedra from NLCE. The shown temperature range corresponds to the part of the curve where the $n = 7$ Euler transform agrees with the $n = 8$ result to ensure convergence. Red crosses indicate the position of the maximum of the specific heat and the inset shows the temperature of the maximum as a function of the applied field h . Bottom: Entropy per site S/N at varying fields as a function of temperature. We show only the converged part of the $n = 8$ Euler transform of our NLCE results. Legends for magnetic field h apply for both panels.




FIG. 9. Magnetization m_z/N as a function of T . We show results obtained from the $n = 8$ Euler transform of the NLCE including data from clusters of up to $n = 8$ tetrahedra, in the range where they agree with the $n = 7$ Euler transform. At low fields $h = 0.4$, $h = 0.8$, we observe a local maximum of the finite temperature magnetization (inset). The dashed lines correspond to the magnetization per site for a single tetrahedron.




FIG. 10. Convergence comparison of different NLCE expansions. Left: Expansion using clusters with complete tetrahedra up to eighth order. Center: Expansion using complete unit cells up to sixth order (100 clusters with $n = 6$ unit cells: included 48 clusters completely with ED and 8 clusters in combination with canonical typicality). Right: Expansion using single sites, up to 14th order. For all expansions the highest order Euler series acceleration is shown. The expansion based on complete tetrahedra consistently yields superior convergence and is used in this work. Dashed lines show the height and position of the Schottky anomaly as extracted from the left panel for comparison. The red dashed lines correspond to the eighth order Euler transform of the tetrahedra expansion, showing that all expansions agree with this result in the converged regime.

results agree and yield a Curie law. At intermediate temperatures, a difference due to the finite size of the tetrahedron is noticeable, with the selection of different total magnetization ground states in the low- T limit, with m_z/N being 0, 0.25, and 0.5, respectively, depending on the field [92].

Interestingly, a nonmonotonic dependence of the magnetization on T can be observed at low fields, Fig. 9, both in the NLCE results (emphasized in the inset) as well as—more visibly—in the single tetrahedron case: the magnetization vanishes at both low T , because the weak-field low-energy states are dominantly nonmagnetic; and at high T , for the usual entropic reasons. At intermediate T , by contrast, the large weight of magnetic states oriented by the fields dominates, hence the maximum.

V. DISCUSSION AND EXPERIMENT

Having laid out the results, we now place them in the broader context of other highly frustrated model systems on one hand, and experiments on magnetic materials on the other. Here we focus on the specific heat, not only because it is a quantity which is readily available across the board; but also because it allows for a relatively straightforward comparison between different platforms thanks to the integral $\frac{1}{N} \int_0^\infty dT' C_V(T')/T' = \ln(2S + 1)$, which is the total entropy of a spin S system. This relation allows for a direct comparison of different S specific heat curves, by an appropriate rescaling to remove the different factors $\ln(2S + 1)$.

We have collated the data for a number of models and materials in Fig. 1. There we compare (i) our converged results with what is found for single tetrahedra with (ii) $S = \frac{1}{2}$ or (iii) in the classical limit $S = \infty$; as well as experiments on




FIG. 11. Convergence comparison of different NLCE expansions for the susceptibility. Left: Expansion using clusters with complete tetrahedra up to eighth order. Center: Expansion using complete unit cells up to sixth order (100 clusters with $n = 6$ unit cells: included 48 clusters completely with ED). Right: Expansion using single sites, up to 14th order. For all expansions the highest order Euler series acceleration is shown. The expansion based on complete tetrahedra consistently yields superior convergence and is used in the remainder of this work. Dashed lines show the height and position of the Schottky anomaly as extracted from the left panel for comparison.

(iv) the Ising spin ice pyrochlore magnet $\text{Dy}_2\text{Ti}_2\text{O}_7$ and (v) the $S = 1$ Heisenberg antiferromagnet $\text{NaCaNi}_2\text{F}_7$, rescaled by $\ln 2 / \ln 3$ to take into account the greater maximal entropy per site for $S = 1$ compared to $S = \frac{1}{2}$. The T axes of the experimental data, (iv) and (v), have been scaled so that the temperature of their respective maxima coincide with the one of our data; while the single tetrahedron results (ii) and (iii) were scaled to agree in the asymptotic limit of high T .

All of these have in common a considerable spectral weight downshift—at $T = 0.25$, all of them exhibit a significant residual entropy, see inset of Fig. 1. There are, however, considerable differences of detail [leaving aside case (iii) on account of the unbounded classical entropy]. The single $S = \frac{1}{2}$ tetrahedron, (ii), releases its entropy more swiftly at low T than our NLCE results on account of its singlet gap coupled with a small residual entropy $S(0) = \frac{1}{4} \ln 2$. The spin ice experiment (iv), with a slightly higher residual entropy of around $S_p = \frac{1}{2} \ln \frac{3}{2}$, in fact releases its entropy even more swiftly, with the peak in the specific heat peak being the most narrow on the low- T side.

The $\text{NaCaNi}_2\text{F}_7$ experiment (v) shows an initial high- T release of the entropy remarkably close to that of our $S = \frac{1}{2}$ results. However, already above the peak in C_V , the release in $\text{NaCaNi}_2\text{F}_7$ is comparatively considerably greater, meaning that the spectral weight downshift in our results is stronger.

Indeed, the breadth of the peak in C_V we find for the $S = \frac{1}{2}$ Heisenberg model is broader not only than all the cases (ii)–(v), but also than the other paradigmatic highly frustrated $S = \frac{1}{2}$ Heisenberg model, that on the kagome lattice. By comparison with the high order series expansion results in Ref. [93], we find that in the isotropic $S = \frac{1}{2}$ Heisenberg antiferromagnet on the kagome lattice, the residual entropy of $0.475 \ln 2$ is reached already at a higher temperature of $T \approx 0.30$, whereas in the pyrochlore lattice our results suggest that

TABLE II. Number of connected $|\mathcal{C}_n|$, symmetrically distinct $|\mathcal{S}_n|$, and topologically distinct clusters $|\mathcal{T}_n|$ per unit cell listed with the order of expansion n (equals system size) for the single site expansions.

n	$ \mathcal{C}_n $	$ \mathcal{S}_n $	$ \mathcal{T}_n $
1	4	2	1
2	12	2	1
3	44	8	2
4	182	19	3
5	816	84	5
6	3856	338	10
7	18916	1650	19
8	95436	8026	41
9	492124	41370	88
10	2582256	215564	207
11	13743828	1147137	483
12	74022676	6170524	1216
13	402692008	33567270	3049
14	2209562820	184140685	8002

this entropy is retained at a lower temperature of $T \approx 0.254$, corresponding to the larger spectral downshift in pyrochlore.

It should be emphasized that this is not at all what would obviously have been expected. Generally, low dimensionality is considered to favor spectral weight downshift, as encoded, e.g., by the Mermin-Wagner theorem. Also, in the Ising setting, triangular motifs are considerably more frustrated— $S_{\text{kagome}} = \frac{1}{24\pi^2} \int_0^{2\pi} \int_0^{2\pi} dx dy \ln \{21 - 4[\cos x + \cos y + \cos(x+y)]\} \approx 0.50183$ [94] for the kagome Ising magnet, *much* larger than in the pyrochlore case $S_{\text{Pauling}} \approx \frac{1}{2} \ln \frac{3}{2} \approx 0.2027$ [4,95].

This implies that there is huge scope for unusual behavior of this model at low T . Alas, our results provide little indication of the detailed nature of the low-energy space of states. Indeed, many proposals for the behavior of this magnet have been made, and it is hard to choose between them based on presently available information, as there is not even compelling evidence in favor of a particular physical picture. The concurrent lack of a pristine experimental realization goes a long way towards explaining the divergence

 TABLE III. Number of connected $|\mathcal{C}_n|$, symmetrically distinct $|\mathcal{S}_n|$, and topologically distinct clusters $|\mathcal{T}_n|$ per unit cell listed with the order of expansion n and system sizes N for the unit cell expansions.

n	N	$ \mathcal{C}_n $	$ \mathcal{S}_n $	$ \mathcal{T}_n $
0	1	4	2	1
1	4	1	1	1
2	8	6	2	1
3	12	50	12	3
4	16	475	90	8
5	20	4881	844	25
6	24	52835	8912	100
7	28	593382	99252	466
8	32	6849415	1142759	2473


 TABLE IV. Number of connected $|\mathcal{C}_n|$, symmetrically distinct $|\mathcal{S}_n|$, and topologically distinct clusters $|\mathcal{T}_n|$ per unit cell listed with the order of expansion n and system sizes N for the tetrahedra expansions.


n	N	$ \mathcal{C}_n $	$ \mathcal{T}_n $
0	1	4	1
1	4	2	1
2	7	4	1
3	10	12	1
4	13	44	2
5	16	182	3
6	18, 19	796	6
7	21, 22	3612	10
8	24, 25	16786	24
9	26, 27, 28	79426	49

of theoretical predictions [15–17,24,89,96] so that different methods arguably come up with the conclusion most suited to them.

We are therefore left with the twin higher-level insights, namely that the pyrochlore $S = \frac{1}{2}$ Heisenberg magnet is at least as frustrated, and arguably interesting, as the one on the kagome lattice; and that it is at least as intractable. We hope that future work will be able to build on the advances reported in this work. And, of course, that the low- T regime will become accessible experimentally in a suitable magnetic material.

Note added: As we were concluding this work, a preprint [97] appeared which also studied the thermodynamic properties of the pyrochlore $S = \frac{1}{2}$ Heisenberg model using a combination of methods including canonical typicality, high temperature series expansion, and the entropy method. It placed particular emphasis on extrapolation schemes in order to access the low- T regime.


 FIG. 12. Number of connected $|\mathcal{C}_n|$ and topologically distinct clusters $|\mathcal{T}_n|$ for the three different expansions.



All eight topologically distinct clusters with one through five tetrahedra.




FIG. 13. All six topologically distinct clusters with six tetrahedra.


ACKNOWLEDGMENTS

We are very grateful to Kemp Plumb and Art Ramirez for kindly supplying the experimental data on $\text{NaCaNi}_2\text{F}_7$ and spin ice. We thank Owen Benton, Ludovic Jaubert, Paul McClarty, Jeffrey Rau, Johannes Richter, Oleg Derzhko, Masafumi Udagawa, and Karlo Penc for very helpful discussions. We acknowledge financial support from the Deutsche Forschungsgemeinschaft through SFB 1143 (project-id 247310070) and cluster of excellence ct.qmat (EXC 2147, project-id 390858490). I.H. was supported in part by the Hungarian National Research, Development and Innovation Office (NKFIH) through Grant No. K120569 and the Hungarian Quantum Technology National Excellence Program (Project No. 2017-1.2.1-NKP-2017-00001). Some of the data presented here were produced using the SYTEN toolkit, originally created by Claudius Hubig [71,72].

APPENDIX A: COMPARISON OF NLCE EXPANSIONS

The NLCE expansion is general in that it permits in principle a wealth of different cluster expansions based on constraints on the included class of clusters. Since the number of clusters grows factorially with the number of constituents, it is often wise to constrain the class sufficiently in order to get a manageable number of clusters at the largest cluster sizes which are still solvable in full diagonalization.

These constraints should respect the underlying physical properties of the system as much as possible to get a rapidly converging series expansion. We show a comparison of three different NLCE expansions for the pyrochlore lattice in Figs. 10 and 11: the single site expansion, the unit cell expansion (i.e., clusters on the fcc lattice, decorated by the tetrahedral unit cell), and the tetrahedral expansion used in the main text.



All 10 topologically distinct clusters with seven tetrahedra.




FIG. 14. All 24 topologically distinct clusters with eight tetrahedra.

All expansions converge to the same curve at high temperatures, however the highest order single site expansion does not reach temperatures low enough to resolve the maximum of the specific heat, even after the Euler transform is applied. The unit cell expansion includes in principle much larger clusters (up to 6 unit cells correspond to 24 site clusters), however it is barely possible to obtain converged results for lower temperatures than in the single site expansion, while the tetrahedral expansion yields convergence to much lower temperatures and a clear resolution of the specific heat capacity maximum.

There are several reasons for this superior behavior: First, the central motif of the pyrochlore lattice is the tetrahedron. It is crucial to avoid dangling bonds and incomplete tetrahedra, which are present in both the unit cell and single site expansion.

Second, due to the construction of unit cell clusters, the clusters used here do not include similarly large loops as in the tetrahedral expansion, which are crucial at low temperatures. Third, and this is purely technical, due to unfavorable symmetry properties of the $n = 6$ unit cell clusters, we were unable to solve all 100 topologically distinct clusters with full diagonalization, since the largest remaining symmetry blocks were too large. Finally, the Euler transform for the unit cell expansion can only rely on five complete expansion orders, which is far less than in the case of the single site and tetrahedral expansion.

Therefore, we conclude that the tetrahedral expansion is far superior to any other approach in the pyrochlore lattice and allows us to access the lowest temperatures.




FIG. 15. 48 site cluster used in our finite temperature DMRG simulations. We use periodic boundary conditions (periodic bonds not shown for clarity).

APPENDIX B: NLCE CLUSTERS

As mentioned in the previous Appendix, a successful NLCE scheme needs to limit the growth of the number of clusters to an extent that the number of largest tractable clusters (using full diagonalization) is not too large.

In Tables II through IV we list the number of clusters appearing at each order n in the single site, unit cell, and tetrahedral expansion. The order n refers to the number of sites, unit cells, or tetrahedra, respectively. The numbers listed correspond to the total number of clusters $|\mathcal{C}_n|$, the number

of clusters which are not identical under application of non-translational symmetries $|\mathcal{S}_n|$, and the number of topologically distinct clusters $|\mathcal{T}_n|$, which is computationally relevant since this is the number of clusters for which the Hamiltonian has to be diagonalized. The growth of the number of clusters with order n is depicted in Fig. 12.

It should be noted that it is computationally challenging to check the topological equivalence of two clusters, since their interaction graphs need to be checked for an isomorphism, which is an NP hard problem. Therefore, e.g., the reduction of 184140685 symmetrically distinct clusters at $n = 14$ in the unit cell clusters to 8002 topologically distinct clusters is already difficult and limits severely the access to higher orders.

For the tetrahedral expansion there is only one topologically distinct cluster for the orders 1 to 3 and two clusters with four tetrahedra. At the highest order we could reach, there are 24 clusters composed of eight tetrahedra, which have either 24 or 25 spins, just at the limit of what can be solved with full exact diagonalization using all symmetries of the clusters. We show all topologically distinct clusters included in the NLCE up to eight tetrahedra in Figs. 13 and 14.

APPENDIX C: FINITE SIZE CLUSTERS

In the present study we consider two finite size clusters, the first being standard 32 site cluster consisting of two unit cells in direction $\vec{a}_1, \vec{a}_2, \vec{a}_3$ studied, e.g., in Refs. [86,97], which we depicted in the inset of Fig. 3. The second cluster we study is the 48 site cluster shown in Fig. 15.

-
- [1] J. S. Gardner, M. J. P. Gingras, and J. E. Greedan, Magnetic pyrochlore oxides, *Rev. Mod. Phys.* **82**, 53 (2010).
 - [2] W. H. Bragg, The structure of magnetite and the spinels, *Nature (London)* **95**, 561 (1915).
 - [3] S.-H. Lee, H. Takagi, D. Louca, M. Matsuda, S. Ji, H. Ueda, Y. Ueda, T. Katsufuji, J.-H. Chung, S. Park, S.-W. Cheong, and C. Broholm, Frustrated magnetism and cooperative phase transitions in spinels, *J. Phys. Soc. Jpn.* **79**, 011004 (2010).
 - [4] P. W. Anderson, Ordering and antiferromagnetism in ferrites, *Phys. Rev.* **102**, 1008 (1956).
 - [5] M. J. Harris, S. T. Bramwell, D. F. McMorrow, T. Zeiske, and K. W. Godfrey, Geometrical Frustration in the Ferromagnetic Pyrochlore $\text{Ho}_2\text{Ti}_2\text{O}_7$, *Phys. Rev. Lett.* **79**, 2554 (1997).
 - [6] C. Castelnovo, R. Moessner, and S. L. Sondhi, Spin ice, fractionalization, and topological order, *Annu. Rev. Condens. Matter Phys.* **3**, 35 (2012).
 - [7] K. W. Plumb, H. J. Changlani, A. Scheie, S. Zhang, J. W. Krizan, J. A. Rodriguez-Rivera, Y. Qiu, B. Winn, R. J. Cava, and C. L. Broholm, Continuum of quantum fluctuations in a three-dimensional $S = 1$ Heisenberg magnet, *Nat. Phys.* **15**, 54 (2019).
 - [8] A. P. Ramirez, A. Hayashi, R. J. Cava, R. Siddharthan, and B. S. Shastry, Zero-point entropy in spin ice, *Nature (London)* **399**, 333 (1999).
 - [9] J. Villain, Insulating spin glasses, *Z. Phys. B: Condens. Matter* **33**, 31 (1979).
 - [10] R. Moessner and J. T. Chalker, Properties of a Classical Spin Liquid: The Heisenberg Pyrochlore Antiferromagnet, *Phys. Rev. Lett.* **80**, 2929 (1998).
 - [11] U. Hizi and C. L. Henley, Anharmonic ground state selection in the pyrochlore antiferromagnet, *Phys. Rev. B* **80**, 014407 (2009).
 - [12] C. L. Henley, Order by Disorder and Gauge-like Degeneracy in a Quantum Pyrochlore Antiferromagnet, *Phys. Rev. Lett.* **96**, 047201 (2006).
 - [13] A. B. Harris, A. J. Berlinsky, and C. Bruder, Ordering by quantum fluctuations in a strongly frustrated Heisenberg antiferromagnet, *J. Appl. Phys.* **69**, 5200 (1991).
 - [14] M. Isoda and S. Mori, Valence-bond crystal and anisotropic excitation spectrum on 3-dimensionally frustrated pyrochlore, *J. Phys. Soc. Jpn.* **67**, 4022 (1998).
 - [15] B. Canals and C. Lacroix, Pyrochlore Antiferromagnet: A Three-Dimensional Quantum Spin Liquid, *Phys. Rev. Lett.* **80**, 2933 (1998).
 - [16] E. Berg, E. Altman, and A. Auerbach, Singlet Excitations in Pyrochlore: A Study of Quantum Frustration, *Phys. Rev. Lett.* **90**, 147204 (2003).
 - [17] R. Moessner, S. L. Sondhi, and M. O. Goerbig, Quantum dimer models and effective Hamiltonians on the pyrochlore lattice, *Phys. Rev. B* **73**, 094430 (2006).
 - [18] F. J. Burnell, S. Chakravarty, and S. L. Sondhi, Monopole flux state on the pyrochlore lattice, *Phys. Rev. B* **79**, 144432 (2009).

- [19] J. Ummethum, J. Schnack, and A. M. Luchli, Large-scale numerical investigations of the antiferromagnetic Heisenberg icosidodecahedron, *J. Magn. Magn. Mater.* **327**, 103 (2013).
- [20] Z. Nussinov, C. D. Batista, B. Normand, and S. A. Trugman, High-dimensional fractionalization and spinon deconfinement in pyrochlore antiferromagnets, *Phys. Rev. B* **75**, 094411 (2007).
- [21] H. Tsunetsugu, Theory of antiferromagnetic Heisenberg spins on a breathing pyrochlore lattice, *Prog. Theor. Exp. Phys.* **2017**, 03310 (2017).
- [22] V. R. Chandra and J. Sahoo, Spin-1/2 Heisenberg antiferromagnet on the pyrochlore lattice: An exact diagonalization study, *Phys. Rev. B* **97**, 144407 (2018).
- [23] A. J. García-Adeva and D. L. Huber, Quantum generalized constant-coupling model for geometrically frustrated antiferromagnets, *Phys. Rev. B* **63**, 174433 (2001).
- [24] Y. Iqbal, T. Müller, P. Ghosh, M. J. P. Gingras, H. O. Jeschke, S. Rachel, J. Reuther, and R. Thomale, Quantum and Classical Phases of the Pyrochlore Heisenberg Model with Competing Interactions, *Phys. Rev. X* **9**, 011005 (2019).
- [25] M. P. Shores, E. A. Nytko, B. M. Bartlett, and D. G. Nocera, A structurally perfect $s = 1/2$ kagomé antiferromagnet, *J. Am. Chem. Soc.* **127**, 13462 (2005).
- [26] J. G. Rau and M. J. P. Gingras, Frustrated quantum rare-earth pyrochlores, *Annu. Rev. Condens. Matter Phys.* **10**, 357 (2019).
- [27] M. J. P. Gingras and P. A. McClarty, Quantum spin ice: A search for gapless quantum spin liquids in pyrochlore magnets, *Rep. Prog. Phys.* **77**, 056501 (2014).
- [28] L. Clark, G. J. Nilsen, E. Kermarrec, G. Ehlers, K. S. Knight, A. Harrison, J. P. Attfield, and B. D. Gaulin, From Spin Glass to Quantum Spin Liquid Ground States in Molybdate Pyrochlores, *Phys. Rev. Lett.* **113**, 117201 (2014).
- [29] Y. Iqbal, T. Müller, K. Riedl, J. Reuther, S. Rachel, R. Valentí, M. J. P. Gingras, R. Thomale, and H. O. Jeschke, Signatures of a gearwheel quantum spin liquid in a spin- $\frac{1}{2}$ pyrochlore molybdate Heisenberg antiferromagnet, *Phys. Rev. Mater.* **1**, 071201 (2017).
- [30] M. Rigol, T. Bryant, and R. R. P. Singh, Numerical Linked-Cluster Approach to Quantum Lattice Models, *Phys. Rev. Lett.* **97**, 187202 (2006).
- [31] M. Rigol, T. Bryant, and R. R. P. Singh, Numerical linked-cluster algorithms. I. Spin systems on square, triangular, and kagomé lattices, *Phys. Rev. E* **75**, 061118 (2007).
- [32] B. Tang, E. Khatami, and M. Rigol, A short introduction to numerical linked-cluster expansions, *Comput. Phys. Commun.* **184**, 557 (2013).
- [33] R. Applegate, N. R. Hayre, R. R. P. Singh, T. Lin, A. G. R. Day, and M. J. P. Gingras, Vindication of $\text{Yb}_2\text{Ti}_2\text{O}_7$ as a Model Exchange Quantum Spin Ice, *Phys. Rev. Lett.* **109**, 097205 (2012).
- [34] R. R. P. Singh and J. Oitmaa, Corrections to Pauling residual entropy and single tetrahedron based approximations for the pyrochlore lattice Ising antiferromagnet, *Phys. Rev. B* **85**, 144414 (2012).
- [35] O. Benton, Instabilities of a U(1) Quantum Spin Liquid in Disordered Non-Kramers Pyrochlores, *Phys. Rev. Lett.* **121**, 037203 (2018).
- [36] O. Benton, L. D. C. Jaubert, R. R. P. Singh, J. Oitmaa, and N. Shannon, Quantum Spin Ice with Frustrated Transverse Exchange: From a π -Flux Phase to a Nematic Quantum Spin Liquid, *Phys. Rev. Lett.* **121**, 067201 (2018).
- [37] L. D. C. Jaubert, O. Benton, J. G. Rau, J. Oitmaa, R. R. P. Singh, N. Shannon, and M. J. P. Gingras, Are Multiphase Competition and Order by Disorder the Keys to Understanding $\text{Yb}_2\text{Ti}_2\text{O}_7$? *Phys. Rev. Lett.* **115**, 267208 (2015).
- [38] N. R. Hayre, K. A. Ross, R. Applegate, T. Lin, R. R. P. Singh, B. D. Gaulin, and M. J. P. Gingras, Thermodynamic properties of $\text{Yb}_2\text{Ti}_2\text{O}_7$ pyrochlore as a function of temperature and magnetic field: Validation of a quantum spin ice exchange Hamiltonian, *Phys. Rev. B* **87**, 184423 (2013).
- [39] A. B. Harris, C. Kallin, and A. J. Berlinsky, Possible Néel orderings of the kagomé antiferromagnet, *Phys. Rev. B* **45**, 2899 (1992).
- [40] J. Röchner, L. Balents, and K. P. Schmidt, Spin liquid and quantum phase transition without symmetry breaking in a frustrated three-dimensional Ising model, *Phys. Rev. B* **94**, 201111(R) (2016).
- [41] H. Tsunetsugu, Antiferromagnetic quantum spins on the pyrochlore lattice, *J. Phys. Soc. Jpn.* **70**, 640 (2001).
- [42] S. V. Isakov, K. Gregor, R. Moessner, and S. L. Sondhi, Dipolar Spin Correlations in Classical Pyrochlore Magnets, *Phys. Rev. Lett.* **93**, 167204 (2004).
- [43] P. Müller, A. Lohmann, J. Richter, and O. Derzhko, Thermodynamics of the pyrochlore-lattice quantum Heisenberg antiferromagnet, *Phys. Rev. B* **100**, 024424 (2019).
- [44] S. Popescu, A. J. Short, and A. Winter, Entanglement and the foundations of statistical mechanics, *Nat. Phys.* **2**, 754 (2006).
- [45] S. Goldstein, J. L. Lebowitz, R. Tumulka, and N. Zanghì, Canonical Typicality, *Phys. Rev. Lett.* **96**, 050403 (2006).
- [46] S. Sugiura and A. Shimizu, Thermal Pure Quantum States at Finite Temperature, *Phys. Rev. Lett.* **108**, 240401 (2012).
- [47] S. Sugiura and A. Shimizu, Canonical Thermal Pure Quantum State, *Phys. Rev. Lett.* **111**, 010401 (2013).
- [48] D. J. Luitz and Y. Bar Lev, The ergodic side of the many-body localization transition, *Ann. Phys. (Berlin)* **529**, 1600350 (2017).
- [49] J. Schnack, J. Richter, T. Heitmann, J. Richter, and R. Steinigeweg, Finite-size scaling of typicality-based estimates, *Z. Naturforsch. Teil A* **75**, 465 (2020).
- [50] T. J. Park and J. C. Light, Unitary quantum time evolution by iterative Lanczos reduction, *J. Chem. Phys.* **85**, 5870 (1986).
- [51] L. A. Knizhnerman, Calculation of functions of unsymmetric matrices using Arnoldi's method, *Zh. Vychisl. Mat. Mat. Fiz.* **31**, 5 (1991).
- [52] J. Jaklič and P. Prelovšek, Lanczos method for the calculation of finite-temperature quantities in correlated systems, *Phys. Rev. B* **49**, 5065 (1994).
- [53] C. Moler and C. Van Loan, Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later, *SIAM Rev.* **45**, 3 (2003).
- [54] I. P. McCulloch, From density-matrix renormalization group to matrix product states, *J. Stat. Mech.: Theory Exp.* **10** (2007) P10014.
- [55] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011), January 2011 Special Issue.
- [56] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, Matrix Product Density Operators: Simulation of Finite-Temperature and Dissipative Systems, *Phys. Rev. Lett.* **93**, 207204 (2004).

- [57] S. R. White, Minimally Entangled Typical Quantum States at Finite Temperature, *Phys. Rev. Lett.* **102**, 190601 (2009).
- [58] E. M. Stoudenmire and S. R. White, Minimally entangled typical thermal state algorithms, *New J. Phys.* **12**, 055026 (2010).
- [59] A. E. Feiguin and S. R. White, Finite-temperature density matrix renormalization using an enlarged Hilbert space, *Phys. Rev. B* **72**, 220401(R) (2005).
- [60] B. Bruognolo, Z. Zhu, S. R. White, and E. M. Stoudenmire, Matrix product state techniques for two-dimensional systems at finite temperature, [arXiv:1705.05578](https://arxiv.org/abs/1705.05578).
- [61] S. Paeckel, T. Khler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, Time-evolution methods for matrix-product states, *Ann. Phys.* **411**, 167998 (2019).
- [62] G. Vidal, Efficient Simulation of One-Dimensional Quantum Many-Body Systems, *Phys. Rev. Lett.* **93**, 040502 (2004).
- [63] A. J. Daley, C. Kollath, U. Schollwöck, and G. Vidal, Time-dependent density-matrix renormalization-group using adaptive effective Hilbert spaces, *J. Stat. Mech.: Theory Exp.* **04** (2004) P04005.
- [64] J. J. García-Ripoll, Time evolution of matrix product states, *New J. Phys.* **8**, 305 (2006).
- [65] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Time-Dependent Variational Principle for Quantum Lattices, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [66] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, Unifying time evolution and optimization with matrix product states, *Phys. Rev. B* **94**, 165116 (2016).
- [67] C. Hubig, I. P. McCulloch, U. Schollwöck, and F. A. Wolf, Strictly single-site DMRG algorithm with subspace expansion, *Phys. Rev. B* **91**, 155115 (2015).
- [68] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [69] S. R. White, Density-matrix algorithms for quantum renormalization groups, *Phys. Rev. B* **48**, 10345 (1993).
- [70] K. A. Hallberg, New trends in density matrix renormalization, *Adv. Phys.* **55**, 477 (2006).
- [71] C. Hubig, F. Lachenmaier, N.-O. Linden, T. Reinhard, L. Stenzel, A. Swoboda, and M. Grundner, The SYTEN toolkit.
- [72] C. Hubig, Symmetry-protected tensor networks, Ph.D. thesis, LMU München, 2017.
- [73] M. Qin, C.-M. Chung, H. Shi, E. Vitali, C. Hubig, U. Schollwöck, S. R. White, and S. Zhang, Absence of superconductivity in the pure two-dimensional Hubbard model, *Phys. Rev. X* **10**, 031016 (2020).
- [74] M. Rigol, T. Bryant, and R. R. P. Singh, Numerical linked-cluster algorithms. II. t - j models on the square lattice, *Phys. Rev. E* **75**, 061119 (2007).
- [75] E. Khatami, J. S. Helton, and M. Rigol, Numerical study of the thermodynamics of clinoatacamite, *Phys. Rev. B* **85**, 064401 (2012).
- [76] T. Pardini, A. Menon, S. P. Hau-Riege, and R. R. P. Singh, Local entanglement and confinement transitions in the random transverse-field Ising model on the pyrochlore lattice, *Phys. Rev. B* **100**, 144437 (2019).
- [77] E. Khatami and M. Rigol, Thermodynamics of strongly interacting fermions in two-dimensional optical lattices, *Phys. Rev. A* **84**, 053611 (2011).
- [78] E. Khatami and M. Rigol, Effect of particle statistics in strongly correlated two-dimensional Hubbard models, *Phys. Rev. A* **86**, 023633 (2012).
- [79] E. Khatami and M. Rigol, Thermodynamics of the antiferromagnetic Heisenberg model on the checkerboard lattice, *Phys. Rev. B* **83**, 134431 (2011).
- [80] E. Khatami, R. R. P. Singh, and M. Rigol, Thermodynamics and phase transitions for the Heisenberg model on the pinwheel distorted kagome lattice, *Phys. Rev. B* **84**, 224411 (2011).
- [81] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, 2nd ed. (Cambridge University Press, Cambridge, 1992).
- [82] R. W. Hamming, *Numerical Methods for Scientists and Engineers*, 2nd ed. (Dover, New York, 1987), p. 721.
- [83] A. W. Sandvik, Computational studies of quantum spin systems, *AIP Conf. Proc.* **1297**, 135 (2010).
- [84] B. D. McKay and A. Piperno, Practical graph isomorphism II, *J. Symb. Comput.* **60**, 94 (2014).
- [85] J. Konc and D. Janež, An improved branch and bound algorithm for the maximum clique problem, *MATCH Commun. Math. Comput. Chem.* **58**, 569 (2007).
- [86] H. J. Changlani, Quantum versus classical effects at zero and finite temperature in the quantum pyrochlore $\text{Yb}_2\text{Ti}_2\text{O}_7$, [arXiv:1710.02234](https://arxiv.org/abs/1710.02234).
- [87] D. J. Klein, Exact ground states for a class of antiferromagnetic Heisenberg models with short-range interactions, *J. Phys. A: Math. Gen.* **15**, 661 (1982).
- [88] J. T. Chayes, L. Chayes, and S. A. Kivelson, Valence bond ground states in a frustrated two-dimensional spin-1/2 Heisenberg antiferromagnet, *Commun. Math. Phys.* **123**, 53 (1989).
- [89] K. S. Raman, R. Moessner, and S. L. Sondhi, $\text{SU}(2)$ -invariant spin-1/2 Hamiltonians with resonating and other valence bond phases, *Phys. Rev. B* **72**, 064413 (2005).
- [90] We hope that there will be no confusion with the bond dimension in DMRG, for which we have conventionally used the same symbol.
- [91] Y. Huang, K. Chen, Y. Deng, N. Prokofev, and B. Svistunov, Spin-Ice State of the Quantum Heisenberg Antiferromagnet on the Pyrochlore Lattice, *Phys. Rev. Lett.* **116**, 177203 (2016).
- [92] Note that accidental degeneracies of the ground state across different S_z sectors lead to intermediate values of the magnetization of the single tetrahedron at $h = 2$ and $h = 4$.
- [93] G. Misguich and B. Bernu, Specific heat of the $S=12$ Heisenberg model on the kagome lattice: High-temperature series expansion analysis, *Phys. Rev. B* **71**, 014417 (2005).
- [94] K. Kan and S. Naya, Antiferromagnetism. The kagome Ising net, *Prog. Theor. Phys.* **10**, 158 (1953).
- [95] L. Pauling, The structure and entropy of ice and of other crystals with some randomness of atomic arrangement, *J. Am. Chem. Soc.* **57**, 2680 (1935).
- [96] J. N. Reimers, A. J. Berlinsky, and A.-C. Shi, Mean-field approach to magnetic ordering in highly frustrated pyrochlores, *Phys. Rev. B* **43**, 865 (1991).
- [97] O. Derzhko, T. Hutak, T. Krokhamalskii, J. Schnack, and J. Richter, Adapting Planck's route to investigate the thermodynamics of the spin-half pyrochlore Heisenberg antiferromagnet, *Phys. Rev. B* **101**, 174426 (2020).