



HAL
open science

Itinerant electrons in the Coulomb phase

L. D. C. Jaubert, Swann Piatecki, Masudul Haque, R. Moessner

► **To cite this version:**

L. D. C. Jaubert, Swann Piatecki, Masudul Haque, R. Moessner. Itinerant electrons in the Coulomb phase. *Physical Review B: Condensed Matter and Materials Physics (1998-2015)*, 2012, 85, pp.054425. 10.1103/PhysRevB.85.054425 . hal-01541717

HAL Id: hal-01541717

<https://hal.science/hal-01541717>

Submitted on 19 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Itinerant electrons in the Coulomb phase

L. D. C. Jaubert,¹ Swann Piatecki,^{1,2} Masudul Haque,¹ and R. Moessner¹

¹Max-Planck-Institut für Physik komplexer Systeme, 01187 Dresden, Germany.

²Laboratoire de Physique Statistique, École Normale Supérieure, UPMC,
Université Paris Diderot, CNRS, 24 rue Lhomond, 75005 Paris, France.

(Dated: October 30, 2013)

We study the interplay between magnetic frustration and itinerant electrons. For example, how does the coupling to mobile charges modify the properties of a spin liquid, and does the underlying frustration favor insulating or conducting states? Supported by Monte Carlo simulations, our goal is in particular to provide an analytical picture of the mechanisms involved. The models under considerations exhibit Coulomb phases in two and three dimensions, where the itinerant electrons are coupled to the localized spins via double exchange interactions. Because of the Hund coupling, magnetic loops naturally emerge from the Coulomb phase and serve as conducting channels for the mobile electrons, leading to doping-dependent rearrangements of the loop ensemble in order to minimize the electronic kinetic energy. At low electron density ρ , the double exchange coupling mainly tends to segment the very long loops winding around the system into smaller ones while it gradually lifts the extensive degeneracy of the Coulomb phase with increasing ρ . For higher doping, the results are strongly lattice dependent, displaying loop crystals with a given loop length for some specific values of ρ , which can melt into another loop crystal by varying ρ . Finally, we contrast this to the qualitatively different behavior of analogous models on kagome or triangular lattices.

I. INTRODUCTION

The combination of magnetism and itinerant electrons is a multi-faceted field in the physics of correlated electrons, where our understanding is still remarkably patchy: even in the case of a square lattice Hubbard model, we lack consensus on a detailed phase diagram in the doping-temperature plane.

Besides the cuprate superconductors, there are plenty of other settings in which interesting questions arise, not least popularized of late by questions raised by the pnictide superconductors, where magnetic frustration and accidental degeneracies have started to be considered.

More broadly, there has been increased interest in the interaction of frustrated magnetism with itinerant electrons^{1–11}. Here, we take up the spirit of this thread of work and study itinerant electrons on a highly frustrated lattice, the pyrochlore lattice. We consider both three and two dimensions, the latter case also being known as the square lattice with crossings, planar pyrochlore, or checkerboard lattice.

We start with an exotic frustrated phase of a magnetic insulator, the Coulomb phase, which has been extensively studied recently^{12–17}. This phase has a number of unusual properties, including algebraic spin correlations and the emergence of extended one-dimensional degrees of freedom^{18–20}, the nature of which is an independently interesting problem²¹.

Especially the latter will play an important role in the following analysis, given the ability of electrons to provide evidence for non-local structures through properties related to transport phenomena. Indeed, it is the marriage of the local constraints imposed by frustration with the ‘non-local’ physics describing mobile particles which makes up for much of the interest in this field.

The Ising-double exchange model which we study here has many parameters: electron density, ρ , temperature T , Ising anisotropy, Hund’s coupling, J_H , magnetic exchange, J , and electron hopping integral, t . A full study of general parameter choices is well-nigh impossible in any detail analytically.

Our approach to the problem considers a regime where the effects of frustration are particularly strong but where considerable progress towards a detailed description is nonetheless possible by analytical (or simple numerical) means. That a (non-trivial) regime where this is possible exists at all is a priori not obvious, and we find that we need to restrict a number of parameters to limiting values: we study the limit where the magnetic energy scales are much larger than the hopping integral t , so that the resulting problem is one of electrons hopping on a classical background spin configuration. We do not need to restrict the electron density to be small, although for that case, we have the most detailed set of results.

We find that the low-density behaviour can be mapped onto a study of a classical loop model with non-trivial weights arising from the addition of electrons Hund’s coupled to the spins in the Coulomb phase. This results in phenomena such as a transition from a (in two dimensions, critical) percolation situation to one in which the loops acquire an exponential length distribution, thereby removing all conducting paths across the sample. As the doping is increased further, we find a sequence of density-dependent preferred loop lengths, which lead to a tendency to form loop crystals which may, however, be frustrated by the lattice geometry.

The organization of this article, and our main results, are summarized below.

A. Summary and Overview

We restrict ourselves to the limit of large exchange coupling interactions so that the ice-rules (see figure 1) themselves are never compromised. We make extensive use of the loop picture encoding the ice rules.²⁰ The loops serve as 1D channels for the electrons. The problem is thus transformed to entropy and energy considerations of possible loop coverings, with loops supporting varying numbers of electron. This “loop framework” for describing conduction electrons is in-

roduced in Section II.

Ref. 20 has described in some detail the loop distributions for both 2D and 3D in the absence of electrons. Once electrons are added, within each loop electrons can occupy states whose energy is given by a 1D dispersion. The dispersion minimum is independent of loop length, so the first electron in a loop has the same energy in all loops. As a result, there is a low doping regime where it is possible to fit at most one electron per loop. All such configurations have the same energy. This is the entropic regime, because entropic arguments determine favorable configurations within an equal-energy manifold.

When the density of electrons ρ is larger, a sub-extensive number of states, possibly even a unique one, tend to be favored, because they manage to minimize the kinetic energy; this we term the energetic regime.

In Section III, we present entropic considerations relevant to the regime of low doping. The total entropy contribution comes from both the loops (with parameters extracted numerically), and the electrons (derived analytically). In particular, the presence of electrons acts like a cutoff on the total number of loops in the system. In 3D, forbidding configurations with less loops than electrons suppresses the formation of extensive loops (present in the peak of the probability distribution function (PDF), see e.g. figure 5), but does not modify the exponents of the PDF. On the other hand in 2D, there is evidence for a variation in the exponent of the power law of the loop distribution.

Once the electronic density gets too large for loops to be restricted to at most one electron, we need to consider energetics. Section IV presents these energy considerations and the phase diagram obtained thereby. Ignoring lattice constraints on loop coverings, we use energy calculations and a Maxwell construction to obtain the phase diagram of Figure 9. We verify some of these results through Monte Carlo simulations. We also present constraints imposed by the lattices under considerations and identify loop crystals arising as the doping is varied.

With periodic boundary conditions, loops spanning the system can be divided into segments connecting “opposite” faces, which we call filaments. Because the transmission of electrons through the system can only occur via these conducting channels, section V is dedicated to their statistics, as a function of doping ρ and dimension. In absence of itinerant electrons, the number of filaments grows linearly with (cubic) system size in 3D but remains constant and of $\mathcal{O}(1)$ in 2D. While these behaviors are qualitatively not modified at low doping, the conducting channels turn out to vanish at intermediate values of ρ .

A separate final section is devoted to an outlook which also contains some words on the behaviour of analogous models on the triangular and kagome lattices, which turn out to exhibit qualitatively different properties both from the pyrochlores and from each other: we find a magnetic conducting solid as well as an insulating cooperative paramagnet.

The numerical component of our work involves Monte Carlo simulations of several types; some details are provided in the Appendix.




FIG. 1: In the Coulomb phase, each frustrated unit (crossed squares in 2D checkerboard, left, and tetrahedra in 3D pyrochlore, right) possesses two up and two down spins, respectively colored in blue and red. Those are the so-called ice-rules or divergence free conditions. Connecting spins of the same color forms a network of loops, as illustrated for the checkerboard. This model is equivalent to the nearest neighbour spin ice model.

II. THE SYSTEM AND THE LOOP FRAMEWORK

A. The model

We will focus primarily on the checkerboard and pyrochlore lattices, which are two lattices where a Coulomb phase can appear (see figure 1). The localized magnetic moments are Ising spins \mathbf{S}_i all parallel to a *global* axis, while itinerant electrons can hop on the lattice sites. The Hamiltonian is

$$\begin{aligned} \mathcal{H} = & J \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j - \sum_{\langle i,j \rangle, \alpha} t (c_{i,\alpha}^\dagger c_{j,\alpha} + c_{j,\alpha}^\dagger c_{i,\alpha}) \\ & - J_h \sum_{i,\alpha,\beta} c_{i,\alpha}^\dagger (\sigma_{\alpha,\beta} \cdot \mathbf{S}_i) c_{i,\beta} \end{aligned} \quad (1)$$

where t is the hopping integral between two neighboring sites, $c_{i,\alpha}^\dagger$ ($c_{i,\alpha}$) are creation (annihilation) operators of itinerant electrons of spin α on site i , and $\sigma_{\alpha,\beta}$ are the Pauli matrices. In order of appearance, the terms in equation (1) are the antiferromagnetic nearest-neighbour exchange between the localized spins incorporating magnetic frustration, the hopping term allowing movement of itinerant electrons, whose spins interact with the localized magnetic moments through ferromagnetic Hund coupling (last term).

In this work, we focus on the limit $t \ll J_h \ll J$. In this limit, the highly degenerate ground state of the frustrated system serves as background for the motion of the electrons. Magnetic excitations (violations of ice rules) are not present in this limit. Electrons can only hop between nearest neighbour spins having the same orientation.

At zero temperature, a Néel or ferromagnetic order would give rise to an insulating or metallic state respectively, but a spin liquid provides a network of conducting paths for itinerant electrons. The present work unveils the geometry of this network, as it is influenced by both the lattice and the minimization of the hopping energy.

B. Emergence of loops

The pyrochlore and checkerboard lattices are made of corner-sharing units with four spins; respectively the tetrahedron and the square with crossings. The antiferromagnetic couplings impose the so-called “ice-rules” with zero magnetization per unit, obtained with two spins pointing up and two spins pointing down²². In absence of itinerant electrons, this ground state is highly degenerate with 6 possible configurations per unit. It corresponds to the 6-vertex model in 2 dimensions²³ and can be mapped onto the nearest neighbor spin ice model in 3 dimensions^{24,25}. These spin systems serve as background for the emergent physics of the so-called *Coulomb phase*¹⁶, a gauge theory where the discrete ice-rules under coarse-grained lead to the emergence of a divergence free flux.

Joining spins of the same orientation in every unit, one obtains loops of up spins and loops of down spins (see Fig. 1). The Coulomb phase can thus be described as an ensemble of possible loop coverings. The resulting loop model possesses two flavors (loops of up and down spins), where every site of the premedial lattice¹⁶ is occupied by two loops, one of each flavor, and every bond is visited by one loop only. In previous work by some of the authors²⁰, a detailed account has been given of the statistics and distribution of these loops, both for the 2D and 3D cases.

In the limit of large J_h , all up (down) electrons are only allowed to hop along an up (down) loop, and are constrained to remain within this loop. This reduces the electron dynamics to be one-dimensional whatever the dimension of the lattice.

The 1D hopping restriction allows us to describe electron dynamics in terms of the dispersion of a 1D tight-binding problem, $E_k = -2t \cos k$, with k the 1D momentum along the loop. The 1D momentum along the loop, $k = 2\pi q/\ell$, is discrete for a loop of finite length ℓ ($q = -\ell/2, \dots, \ell/2 - 1$). An up (down) loop of length ℓ can contain between 0 to ℓ up (down) electrons. Note that double occupancy does not occur in the limit we are considering.

The lowest and highest single-particle levels in the dispersion have energy $\pm 2t$, independent of the loop length (figure 2).

Our system is thus described by

- the number of lattice sites N , equal to $4L^2$ and $16L^3$ for the checkerboard and pyrochlore lattices respectively, where L is the linear number of unit cells;
- the total electron number N_e , and the electron density $\rho = N_e/N$;
- the loop histogram of a given configuration, *i.e.* the number h_i of loops of length ℓ_i , ℓ_i being necessarily an even number on a bipartite lattice;
- the type of lattice, which will among other things determine the smallest possible loop in the system ℓ_{min} (4 for checkerboard and 6 for pyrochlore); the longest possible loop length is always $\sim N/2$.

We define a few additional relevant observables:




FIG. 2: Single-particle energy levels of itinerant electrons confined to loops of length $\ell = 8$ (left) and $\ell = 12$ (right), due to the hopping term in Hamiltonian (1). There are nondegenerate levels at the highest and lowest energies, $\pm 2t$, independent of loop length; the other levels are doubly degenerate.

- the total number of loops in a given configuration $N_\ell = \sum_i h_i$;
- the average loop length for a given configuration $\bar{\ell} = \frac{\sum_i h_i \ell_i}{\sum_i h_i} = \frac{N}{N_\ell}$;
- the statistical average loop length $\langle \ell \rangle$ over all loop configurations;
- the statistical average number of loops $\langle N_\ell \rangle$;
- the number of filaments (see section V).

III. LOW DOPING REGIME

In this section we consider the low doping regime where N_e is small enough to have loop configurations with more loops than electrons, $N_\ell > N_e$. Since the lowest single-electron energy level in any loop is $-2t$, the minimum accessible energy is the same ($-2tN_e$) for all such configurations. Therefore, the ground state manifold consists of all such loop coverings with the same energy $-2tN_e$. The free energy within this manifold is then only determined by entropics. At zero temperature, entropy is understood in the sense that all configurations have an equal probability to occur.

In the first subsection below, we present some entropy calculations, combining loop and electronic contributions to the entropy, and show how this determines the average loop length at nonzero electronic density ρ . In the second subsection we present numerical results on the effect of electrons on the entire loop length distribution (loop PDF). The effect on the loop PDF is a natural way to characterize the influence of electrons in the magnetic Coulomb-phase system.




FIG. 3: Histogram of distribution of the number of loops N_ℓ per configuration in 2D (left) and 3D (right), in the absence of electrons. The distribution is scaled according to the Gaussian expression of equation (2) collapsing all system sizes onto the same curve ($L = 100$ to 600 in 2D and $L = 4$ to 60 in 3D). In 2D, the two Gaussians have the same width and corresponds to even and odd N_ℓ (upper and lower curve). In 3D, finite size corrections are visible for $L = 4$ (red crosses).

A. Entropy

The total entropy consists of loop and electronic contributions. The loop distribution is of course itself affected by the itinerant electrons. However, in the limit of small ρ , the change in loop entropy is small. Below we combine the $\rho = 0$ loop entropy with the finite- ρ electronic entropy to approximate the total entropy at small ρ .

Loop entropy. Figure 3 shows the distribution of the number of loops N_ℓ having length ℓ , in the absence of electrons ($\rho = 0$), obtained from Monte Carlo simulations (Appendix A). The distribution has Gaussian form:

$$P(N_\ell) \sim \frac{1}{\sqrt{N}} \exp \left[-\frac{(N_\ell - \langle N_\ell \rangle)^2}{2\kappa N} \right] \quad (2)$$

where $\kappa_{2d} = 0.0384$ and $\kappa_{3d} = 0.00423$. In 2D, there are *two* Gaussians, corresponding to even and odd N_ℓ .

For a system with short-distance correlations, Gaussian distributions are natural to expect from the Central Limit Theorem, since by dividing the system into small mesoscopic segments the total distribution can be recast into a sum of many random variables. In our case, however, we have a system with algebraic correlations and extended objects (loops), so finding Gaussian distributions is not a priori trivial.

Equation 2 can be expressed in terms of the average loop length $\bar{\ell} = N/N_\ell$ instead of N_ℓ :

$$\begin{aligned} P(\bar{\ell}) \equiv P(N_\ell) \left| \frac{dN_\ell}{d\bar{\ell}} \right| &\sim \frac{\sqrt{N}}{\bar{\ell}^2} \exp \left[-\frac{N}{2\kappa} \left(\frac{1}{\bar{\ell}} - \frac{1}{\langle \ell \rangle} \right)^2 \right] \\ &\sim \frac{\sqrt{N}}{\bar{\ell}^2} \exp \left[-\frac{N}{2\kappa} \left(\frac{\bar{\ell} - \langle \ell \rangle}{\langle \ell \rangle^2} \right)^2 \right] \end{aligned} \quad (3)$$

Here we have used $1/\bar{\ell}\langle \ell \rangle \approx 1/\langle \ell \rangle^2$, which is valid for large N in the region where the Gaussian is appreciable. Thus the

loop contribution to the entropy ($\sim \ln P$) is

$$S_{loop} = S_1 - \frac{N}{2\kappa \langle \ell \rangle^4} (\bar{\ell} - \langle \ell \rangle)^2, \quad (4)$$

where S_1 is a constant.

Electronic entropy. Since we have at most one electron per loop, the number of possible combinations to put N_e electrons in N_ℓ loops is the binomial (Pascal) coefficient $N_\ell!/[N_e!(N_\ell - N_e)!]$. The logarithm then gives the electronic contribution to the entropy. Using Stirling's approximation for the thermodynamic limit ($N_\ell \gg 1$, $N_e \gg 1$), we get as per usual

$$S_{elec} = \frac{N}{\bar{\ell}} [x \ln x + (1-x) \ln(1-x)] \quad (5)$$

where $x = N_e/N_\ell = \rho\bar{\ell}$.

Total Entropy. For small electron densities, we can approximate the loop entropy by the $\rho = 0$ expression calculated above. The total entropy of electrons and loops is then

$$\begin{aligned} S_{tot} = & -\frac{N}{\bar{\ell}} [\rho\bar{\ell} \ln \rho\bar{\ell} + (1-\rho\bar{\ell}) \ln(1-\rho\bar{\ell})] \\ & + S_1 - \frac{N}{2\kappa \langle \ell \rangle_0^4} (\bar{\ell} - \langle \ell \rangle_0)^2. \end{aligned} \quad (6)$$

The subscript 0 represents the ensemble average at $\rho = 0$. The value of $\bar{\ell}$ that maximizes S_{tot} is found (in the limit $\rho\bar{\ell} \ll 1$) to be

$$\bar{\ell}_{opt} \approx \langle \ell \rangle_0 - \rho\kappa \langle \ell \rangle_0^3. \quad (7)$$

For $\rho\langle \ell \rangle_0 \ll 1$, the distribution stays almost Gaussian, so this most probable value of $\bar{\ell}$ is approximately the mean value of the distribution $\langle \ell \rangle_\rho$. In Figure 4, we compare this prediction

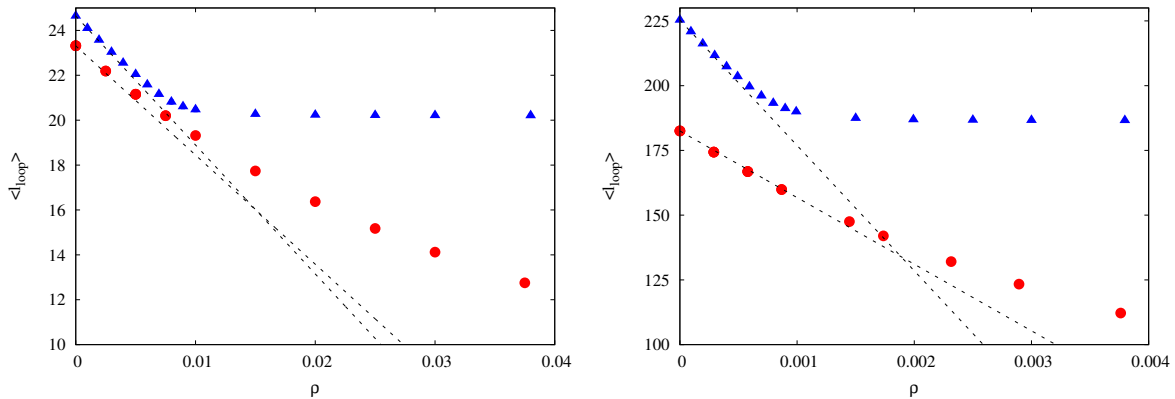


FIG. 4: Mean value of the loop length $\langle \ell \rangle$ as a function of doping. *Left*: checkerboard, $L = 10$ (red dots) and 60 (blue triangles). *Right*: pyrochlore, $L = 6$ (red dots) and 20 (blue triangles). The dashed straight lines are predictions from Equation (7), without any fitting parameters. For both panels, the blue triangles data points are very close to thermodynamic limit.

to numerical data. For each system size L , we can extract $\langle \ell \rangle_0$ from the numerical $\langle \ell \rangle$ at $\rho = 0$. Eq. (7) then gives a linear prediction (dashed straight lines in Figure 4), which works well for small ρ .

B. Loop length distribution

We next examine the effect of electrons on the entire loop probability distribution function (PDF). We denote by τ the power-law exponent, when the PDF has form $\ell^{-\tau}$. Without electrons ($\rho = 0$), the PDF follows $P_{2D} \sim L^2 / \ell^\tau$ in 2D with $\tau = 2 + 1/7$. In 3D, the $\rho = 0$ PDF displays a crossover around $\ell \sim L^2$ between two power laws, from $L^3 / \ell^{5/2}$ to $1/\ell$. (For details, see Ref. 20.)

At finite electron densities, loop configurations with $N_\ell < N_e$ are rejected due to energetics, as explained previously. Thus electron doping acts not entirely unlike a “chemical potential” for loops, favoring configurations with more loops, and thus a priori shorter ones. In both 2D and 3D, this implies a disappearance of longer loops, as can be seen in the calculated PDF’s of Figures 5 and 6, where the large- ℓ parts of the PDF are progressively decimated for increasing ρ . In 3D the form of the PDF is otherwise unchanged (Figure 5). In 2D the effect seems to be more drastic; the inset of figure 6 suggests that the entire power-law behavior of the PDF is modified.

To quantify how the 2D loop PDF changes qualitatively at finite ρ , we define and compute a *local* exponent in ℓ : $\tau_{\text{ocal}}(\ell, \rho) = \log(P_{2D}(\ell, \rho) / P_{2D}(2\ell, \rho)) / \log 2$. This is displayed in the upper right inset to Figure 6. Our results suggest a trend toward increasing τ_{ocal} as a function of ρ , consistent with the idea that itinerant electrons favor small loops.

This outcome deserves a few comments. In Ref. 20, the loop statistics of the 2D Coulomb phase (zero doping) has been shown to be analogous to the Stochastic Loewner evolution process $\text{SLE}_{\kappa=6}$ with fractal dimension $D_f = 1 + \kappa/8 = 7/4$. The SLE_κ can be identified to various realizations of the $\mathcal{O}(n)$ model through the relation $n = -2 \cos(4\pi/\kappa)$.²⁶ The $\mathcal{O}(n)$ model is often used to describe fully packed loop




FIG. 5: Loop length distribution $P_{3D}(\ell, \rho)$ in 3D for different values of doping ($\rho = \{0, 0.00075, 0.0015, 0.002\}$) and system size $L = 10$. The distribution is normalized such that $\int \ell P_{3D}(\ell, \rho) d\ell = 16L^3 = N$. The dashed line indicates the power law fit ℓ^{-1} at $\rho = 0$ and the arrow shows the shifting of the distribution for increasing ρ . The exponents of the two power law regions (before and after $\ell \approx L^2 = 100$) do not vary, but the peak for long winding loops $\ell \sim 8L^3 = N/2$ gets smaller with increasing ρ .

models with loop fugacity n (see *e.g.* Ref. 23). The partition function of the fully packed loop model is $\mathcal{Z} = \sum n^{N_\ell}$, where the sum runs over all possible configurations. The Coulomb phase corresponds to a fully packed loop model²⁰; at zero doping, the free energy of our model is trivially independent of the number of loops N_ℓ and thus corresponds to a fugacity $n = 1$. Higher values of the fugacities favor configurations with more loops and tend to increase the value of τ ²⁷, in a way reminiscent of the influence of doping here. The addition of itinerant electrons remains a non-trivial problem and is not exactly the same as a fugacity for a loop, but at small and finite doping ρ , some features could be captured by $\mathcal{O}(n(\rho) > 1)$ models or $\text{SLE}_{\kappa(\rho) < 6}$ processes.




FIG. 6: *Main*: Loop length distribution $P_{2D}(\ell, \rho)$ in 2D for different values of doping ($\rho = \{0, 0.0025, 0.0056, 0.0087\}$) and system size $L = 40$. The distribution is normalized such that $\int \ell P_{2D}(\ell, \rho) d\ell = 4L^2 = N$. The dashed line indicates power law $L^2/\ell^{15/7}$ and the arrow shows the shifting of the distribution for increasing ρ . On this scale, it is not obvious whether the exponent of the power law varies or not. *Inset*: The power-law exponent defined locally on ℓ , τ_{local} for $\ell = 8$ (■), 16 (▲), 32 (●) for $L = 40$ (similar behavior obtained for $L = 20$ and 60).

IV. LARGE DENSITIES; PHASE DIAGRAMS

In this section, we present results relevant for higher densities, where entropic considerations are no longer sufficient, and non-trivial electronic hopping energies need to be considered. We continue to describe the system in terms of loop coverings.

We first provide an analysis based on calculations for loop coverings of equal-length loops. Using a Maxwell construction, we can use this information to predict ranges of electron density where the ground state manifold consists of coverings by loops of two different lengths.

These considerations, described in the first two subsections below, do not take into account any lattice constraints other than the fact that the minimum loop length is $\ell_{\text{min}} = 4(6)$ for the checkerboard (pyrochlore). Lattice constraints, disallowing some loop coverings, are difficult to enumerate or list comprehensively on account of their non-local nature. In the final subsection, we present the lattice constraints that we have identified, and their implications.

A. Equal-length loop configurations

In this subsection and the next, we imagine that coverings with any unique loop length $\ell \geq \ell_{\text{min}}$ are possible. Later, we will show that at certain fillings, the ground states exhibit a unique loop length, while at others, the behaviour can be understood through a Maxwell construction based on the single-length results.

For an odd number of electrons $\rho\ell = 2n_o + 1$ in a loop of

length ℓ , the energy is (see figure 2)

$$E(\rho, \ell) = -2t \sum_{n=-n_o}^{n_o} \cos\left(\frac{2\pi n}{\ell}\right) = -2t \frac{\sin(\pi\rho)}{\sin(\pi/\ell)} \quad (8)$$

For an even number of electrons $2n_o$, this expression becomes

$$E(\rho, \ell) = -2t \frac{\sin \pi (\rho - 1/\ell)}{\sin(\pi/\ell)} - 2t \cos(\pi\rho) \quad (9)$$

More generally, if $\rho\ell$ is not an integer, we define the highest odd integer below $\rho\ell$ as

$$\eta = 2E\left(\frac{\rho\ell - 1}{2}\right) + 1 \quad (10)$$

where $E(\cdot)$ is the floor function. Each loop is filled with at least η electrons up to the energy levels at $k = \pm\pi(\eta - 1)/\ell$, while the remaining $N(\rho - \eta/\ell)$ electrons in the system are distributed in the partially filled level at $k = \pm\pi(\eta + 1)/\ell$. The total energy is

$$E(\rho, \ell) = -2tN \left[\frac{\sin(\pi\eta/\ell)}{\ell \sin(\pi/\ell)} + \left(\rho - \frac{\eta}{\ell}\right) \cos\left(\pi \frac{\eta + 1}{\ell}\right) \right] \quad (11)$$

These expressions are electron-hole symmetric, *i.e.* invariant under $\rho \leftrightarrow (1 - \rho)$.

Figure 7 displays the loop length $\ell(\rho)$ that minimizes the energy (11) as a function of ρ , and thus corresponds to the ground state if we impose a unique loop length in the system. (The electron-hole symmetry shows up clearly through the mirror symmetry on either side of $\rho = 0.5$. Therefore, we shall from now on only consider densities below $1/2$.)

Since the lowest possible energy for an electron is $-2t$ and is only accessible for one electron per loop, loops of length 4 are favored up to $\rho = 1/4$. Of course for lower densities than $1/4$, other configurations may be possible as long as the number of loops is larger than the number of electrons, but a system with only loops of length 4 will always be part of the ground state manifold for $\rho \leq 1/4$.

Remarkably, for most densities, a finite loop length with discrete energy levels is preferred compared to infinite loops, except for $\rho = 1/3$ and $2/5$.

We plot on figure 8 the minimum energy corresponding to the loop length $\ell(\rho)$ of figure 7: $E_{\text{min}}(\rho) = E(\rho, \ell(\rho))$.

B. Maxwell construction; lattice-independent phase diagram

We now move beyond configurations with unique loop length. We need to consider mixtures of electron densities in the different loops. This is done through a Maxwell construction, similar to the physics of a liquid-gas first-order phase transition.

A system of N sites and density ρ can be divided into two subsets of sites N_1 and N_2 , with density ρ_1 and ρ_2 of electrons and loop length ℓ_1 and ℓ_2 respectively, with

$$\begin{aligned} N &= N_1 + N_2 \\ N\rho &= N_1\rho_1 + N_2\rho_2. \end{aligned} \quad (12)$$




FIG. 7: Loop length $\ell(\rho)$ minimizing the energy (11) for a system filled with loops of length ℓ only, assuming no lattice constraints other than $\ell_{min} = 4$ (left panel, checkerboard) or 6 (right panel, pyrochlore). ρ takes all rational values p/q , with $q \in \{0, 1, 2, \dots, 1000\}$ and $p \in \{0, 1, 2, \dots, q\}$, and we consider all loop lengths from ℓ_{min} to $\ell_{max} = 100$ (red crosses) and 1000 (blue squares): plotting two different values of ℓ_{max} provides a graphical way to visualize the two values of ρ where the most favorable loop length is infinite, namely $\rho = 1/3$ and $2/5$ (and their symmetric images with respect to $\rho = 1/2$). If several loop lengths give the same energy, we plot the smallest one.

Now if a straight line between $E_{min}(\rho_1)$ and $E_{min}(\rho_2)$ remains below the curve $E_{min}(\rho)$ on figure 8, *i.e.* if

$$N E_{min}(\rho) > N_1 E_{min}(\rho_1) + N_2 E_{min}(\rho_2), \quad (13)$$

then the mixture of two densities is more stable than a unique density and “phase separation” occurs. We thus construct a phase diagram, separating regions of different loop length combinations.

On figure 9, the dots correspond to densities where a unique loop length is favored, whereas the zones between them are “phase mixed”, defined by the loop lengths and electron densities of the surrounding dots. The ratio of one phase compared to the other is given by equations (12). For example for $\rho = 0.4$, 80% of the sites belong to loops of length $\ell_1 = 8$ with electron density $\rho_1 = 3/8$, while the remaining 20% belong to loops of length $\ell_2 = 6$ with density $\rho_2 = 1/2$. As shown on figure 7, loops of length 4 are particularly robust over a wide range of ρ ; preventing their formation, *e.g.* in the pyrochlore lattice in $d = 3$ which permits loops of minimal length 6, thus strongly modifies the phase diagram.

These results indicate that itinerant electrons tend to favor relatively small loops and to prevent the formation of infinite ones at zero temperature.

At low electron densities (up to $1/\ell_{min}$), the minimum energy for a given ρ is degenerate: *e.g.* for $\rho = 1/12$, having only loops of length 4, 6, 8, 10 or 12 gives the same energy, as the number of electrons is smaller than the number of loops on the system, and every electron can fill the lowest energy level. This explains why $E_{min}(\rho)$ is a straight line in this region (Figure 8): as ρ decreases from $1/\ell_{min}$ to 0, the degeneracy of the ground state increases until one recovers naturally the full degeneracy for $\rho = 0$. This of course corresponds to the entropic regime mentioned above.

We stress again that these results have been obtained independently of the lattice (except for the value of ℓ_{min}), and

that lattice constraints (next subsection) will modify some of the phase diagram.

C. Lattice constraints on loop coverings

We now consider effects of the lattice in disallowing some of the configurations predicted by our analysis above. For the specific densities where it is possible to cover the lattice with loops of a unique length, the result is an ordered *loop crystal*. In some cases, however, a covering by a unique length or by a combination of loops of two lengths, is not possible. We point out some such cases below.

1. Checkerboard

We first focus on the regime $\rho \in [1/4; 3/8]$ favoring loops of length 4 or 8 according to the Maxwell construction. As illustrated in figure 10, not only can the lattice be covered by these loops, but the transformation from 4 loops of length 4 to 2 of length 8 is also purely local and allows all possible ratios between these two phases in the thermodynamic limit. A pair of loops of length 8 (one made of up spins, the other of down ones) cannot be separated if there are no other lengths than 4 and 8 in the system. We shall call such pair a defect. The defect concentration is determined by ρ . Defects are not topological in the sense that they can be created and annihilated locally, they can be placed anywhere in the background made of loops of length 4.

At quarter filling, there is one electron per loop of length 4, its energy is $-2t$ and their total number is $N/4$. Each additional pair of electrons fills the first excited level of the newly created defect (3 electrons per $\ell = 8$ loop) and gives an energy $-4t(\sqrt{2} - 1)$. At a given ρ , the number of elec-




FIG. 8: Minimum energy $E(\rho)$ corresponding to the value of $\ell(\rho)$ plotted on figure 7, for a system filled with loops of length ℓ only, assuming no lattice constraints other than $\ell_{min} = 4$ (left, checkerboard) or 6 (right, pyrochlore). ρ takes all rational values p/q , with $q \in \{0, 1, 2, \dots, 1000\}$ and $p \in \{0, 1, 2, \dots, q\}$. We consider all loop length from ℓ_{min} to $\ell_{max} = 1000$. We chose arbitrarily $t = 1/2$. A Maxwell construction can be visualized from this figure. The red dots are a set of points, such that the dashed lines connected them are always below the red curve $E(\rho)$. Hence a mixture of two phases corresponding to two consecutive red dots (ℓ_1, ρ_1) and (ℓ_2, ρ_2) has a lower energy than a single phase with a unique density of electrons ρ on a unique type of loop length ℓ . For increasing values of ρ , the red dots correspond to the most favored loop length $\ell = \{4, 8, 6, 8, 4\}$ (left panel) and $\ell = \{6, 12, 10, 8, 6, 8, 10, 12, 6\}$ (right panel), as can be read from figure 7.




FIG. 9: Zero temperature “phase diagram” of loop length configuration as a function of electron density ρ , obtained without accounting for lattice constraints other than $\ell_{min}=4(6)$ for checkerboard (pyrochlore). A phase is defined by the length of its constituting loops and the density of electrons on them. The values of ρ and $\ell(\rho)$ for each phase are given below and above the lines.




FIG. 10: Schematic representation of the checkerboard lattice with 16 sites and different loop coverings. *Left*: 4 loops of length 4, corresponding to the ground state at $\rho = 1/4$ with 1 electron per loop. *Center*: 2 loops of length 8, corresponding to the ground state at $\rho = 3/8$ with 3 electrons per loop. Between $1/4$ and $3/8$, a mixture of these two configurations will occur, their respective ratio being set by the total number of electrons $N_e = \rho N$. This arrangement of two loops of length 8 appears as a “defect” in a crystal of loops of length 4. *Right*: a red loop of length 6 spanned by another blue loop; the shortest way to close this blue loop requires 10 sites. It is thus impossible to have a mixture of loops of length 6 and 8 only.

trons added with respect to the loop crystal at quarter filling is $(\rho N - N/4) = (4\rho - 1)L^2$. The total energy of the system

between $\rho = 1/4$ and $3/8$ is then

$$E(\rho) = -2t L^2 \left(1 + (4\rho - 1)(\sqrt{2} - 1) \right) \quad (14)$$

Once normalized per number of sites $4L^2$, this expression is the dashed line plotted on figure 8. We have verified this analytical result with finite-temperature Monte Carlo simulations; some details of the method are in Appendix A. Extensive degeneracy is recovered in this region between $1/4$ and $3/8$.

At higher doping, above $\rho = 3/8$, the Maxwell argument predicts a mixture of loops of length 8 and 6. However, as we can see in figure 10, a single loop of length 6 imposes the presence of loops of length 10 at least. This implies that this region of the phase diagram (obtained without accounting for such lattice constraints) is further modified, in an as yet unknown manner.

2. Pyrochlore

Analogous modification of the phase diagram of figure 9 is more severe for the pyrochlore, and shows up already at $\rho = 1/\ell_{min} = 1/6$. We have found that it is impossible to cover the pyrochlore lattice with loops of length 6 only. However, it is possible to do it for half of the system, as explained next.

The pyrochlore lattice can be seen as an alternative stack of kagome and triangular layers orthogonal to one of the global [111] axes. As depicted on figure 13(d), 2/3 of each kagome layer can be filled with loops of length 6 (blue hexagons), while the other 1/3 of the kagome sites forms extensive winding loops along the [111]-axis, crossing alternatively the kagome and triangular layers. Since all blue sites form loops of length 6, putting one electron per loop provides a ground state configuration up to $\rho = 1/12$ at least. This is a priori not the only one, but this proves its existence.

Indeed, despite an intensive search by complete enumeration of configurations respecting the ice-rules on the pyrochlore lattice, we have not detected any relevant loop crystal or mixture of them for systems up to 128 sites. A system of size 192 has been partially investigated, with the same outcome. The smallest occurrence of a single-length loop covering is for 8 loops of length 16 in a system of 128 sites; however this is not a relevant length according to our phase diagram of Figure 9. Thus, the loop coverings on an actual pyrochlore are modified from Figure 9 in most or all of the density range $\rho \in [1/6, 1/2]$. The details of this modification remains an important open problem.

V. CONDUCTING CHANNELS OR “FILAMENTS”

In this section, we discuss the effect of the Coulomb phase and loop structure on the conductivity of mobile electrons. Since the electrons are confined to loops, they can conduct only if a loop connects one edge of the sample to the opposite end. Therefore, we study the number of such sample-spanning loops, which, following Ref. 28, we refer to as “filaments”.

When there are no filaments (*e.g.* in loop crystals with only finite-length loops), the system is unambiguously an insulator. When there are filaments spanning the system, the system cannot be immediately called a conductor, because the actual conductivity will depend on scattering mechanisms exterior to our model.

In the first subsection below, we consider small dopings, where we present Monte Carlo results for the average number of filaments as a function of system size. In the second subsection, we comment on the consequences of our phase diagram.

A. Filaments at low densities

We first consider very low densities that the loop distribution can be assumed to be largely unchanged from the $\rho = 0$ case.




FIG. 11: Average number of filaments, *i.e.* those segments of loops spanning the entire system from one border to the opposite one, as a function of system size L . We use two definitions of filaments; they can either cross the orthogonal borders thanks to the periodic boundary conditions (●) or not (■). The cartoon on the right shows such a boundary-crossing filament, which would be excluded in the ■ data. Dashed lines are guides to the eye for the linear behavior with L . Both x- and y-axes are on a logarithmic scale.

In 2D, there is a small but constant number (≈ 1.86) of winding loops on average in the thermodynamic limit²⁰. There is thus some probability to have filaments in the checkerboard case, but the number of filaments does not grow with system size.

In 3D, the background of winding loops ensures that there are filaments whose number grows with system size. The data in Figure 11 shows that the number of filaments increases linearly with the linear size of the cubic sample.

Since our simulations and loop counting are performed with periodic boundary conditions, some of the filaments, while spanning the sample in one direction, also cross one of the orthogonal boundaries. (An example is shown in the cartoon to the right of Figure 11.) One can argue that this type of filament would not contribute to conduction in a real-life cubic pyrochlore sample. Therefore, we show data both excluding and including this type of filament, and they are seen to have the same power-law behavior.

Since the data including these filaments have better statistics, for finite ρ we display the inclusive data, with the expectation that there is no qualitative difference.

According to Fig. 5, itinerant electrons tend to make disappear extensive loops of length L^3 in 3D. Whether or not this prevents the formation of filaments is not as straightforward as it seems. Indeed, one could naively assume that the number of conducting channels will decrease and maybe even vanish in the thermodynamic limit. However the number of filaments remains approximately constant as plotted on figure 12. This means that the average number of loops increases with doping via dividing the very long ones into smaller but nonetheless extensive loops spanning the system. In the low doping regime in 3D, the number of conducting channels is approximately *independent* of ρ .

By contrast, in 2D, with relatively large error-bars, the number of filaments decreases but remains of $\mathcal{O}(1)$ as electrons are

added.

B. Conduction channels at larger fillings

In Section IV, we identified densities (for both 2D and 3D) where the system is a loop crystal or is covered by loops of two finite lengths only. In such cases, there are no filaments, and the system is truly insulating: lifting the frustration-induced degeneracy removes the non-locality of the loops. How the energetics (and resulting degeneracies) imposed by lattice constraints (see section IV C) manifest themselves in transport properties is an intriguing separate question.

For some ranges of ρ , especially in 3D, the loop coverings indicated by the Maxwell construction are disallowed by the lattice geometry (section IV C). In such cases it remains an open question whether or not the lattice constraints result in loop coverings including infinite (sample-spanning) loops. Unfortunately, the issue of conducting channels depends on the answer to this generally unresolved question.

VI. SUMMARY AND OUTLOOK

We have analyzed the double exchange model on the pyrochlore lattice in two and three dimensions. We have chosen to consider a parameter range for which (i) magnetic frustration is known to give rise to unconventional ground state ensembles and (ii) where it is possible to make considerable analytical progress by mapping the system onto an ensemble of loops, the statistical properties of which are influenced by the addition of the electrons.

We have identified a number of phenomena which depend on features such as dimensionality, which determines whether or not there exist loop segments winding around the system; or lattice structure, which may frustrate the geometric packing of preferred loop lengths.

The model studied here leaves unanswered a number of questions and immediately suggests many generalizations and extensions. We have worked in a limit of parameters such that the magnetic exchange and anisotropy dominate over the Hund coupling which in turn dominates the hopping integral. Our analysis applies to zero temperature.

It would be interesting to relax any of these choices, although technically this may not be easy. In particular, given the presence of gapless excitations on long loops, interesting low-temperature physics may appear. Canting can give rise to non-trivial Berry phase physics, and finite Hund's coupling will enable electrons to hop between loops. For example, recent work at quarter filling on the pyrochlore Kondo lattice has shown the emergence of a chiral magnetic order in the weak-coupling regime²⁹.

Finally, even in the parameter range discussed here, it will be interesting to ask how different frustrated lattices shape up compared to the pyrochlores. We devote the final paragraphs of this paper to discussing this question for the case of triangle-based lattices, the triangular Bravais lattice and the Archimedean kagome lattice.

A. Other lattices: triangular and kagome

The Ising ground states observed for the pyrochlore lattice have vanishing total spin on each tetrahedron. These 'ice rules' ensure the existence of the Coulomb phase and states obeying them amount to a moderate zero-point entropy of less than a third of that of a free spin. By contrast, the zero-point entropy of the triangular Ising magnet is not far from half of that of a free spin, while that of the kagome magnet is over 70% of $\log 2$.

Most fundamentally, the single triangle is relatively much more degenerate than a tetrahedron, with 6 out of 8 (rather than out of 16) states being ground states. The triangle states have varying magnetizations of ± 1 (whereas states obeying the ice rule have a unique magnetization, unless one tunes a field to a transition between magnetization plateaux³⁰).

To move towards the full triangle-based lattices, it is worth noting that for any Ising antiferromagnet, the 'hopping network' formed by neighbouring aligned spins can have a coordination of at most half that of z , the coordination of the underlying lattice – otherwise it would be energetically favorable to flip the highly coordinated spin. The concept of the hopping network generalizes the loops on which the electrons hop on the planar and three dimensional pyrochlores.

It now turns out that triangle and kagome lattices behave entirely differently from pyrochlore in both two and three dimensions, on account of the nature of their frustrated ground state ensembles.

1. Triangular lattice

The ground-state degeneracy of the triangular lattice is immediately lifted by the addition of even a single hole. This result is entirely analogous to the frustrated Nagaoka theorem presented in Ref. 31, in the context of the magnetic supersolid discussed there, and it is also connected to the triangular Bosonic supersolids^{32–35}.

As $z = 6$ for the triangular lattice, the hopping network of aligned spins no longer has coordination two as in the loops of the pyrochlore lattice. In fact, the coordination of a site of the hopping network no longer even needs to be uniform, so that there may be dangling or even isolated sites, as on Fig. 13.(a). The coordination can range all the way from zero (for a spin surrounded by a hexagon of oppositely aligned spins) to maximally $z/2 = 3$.

The latter happens when the hole sits on a site experiencing zero net exchange field from its six neighbors. There exists a unique state (pictured in Fig. 13.(b)) in which there is a network of three-fold coordinated sites. This state breaks translational symmetry by tripling the unit cell as well as time-reversal symmetry as it corresponds to a state with a magnetization of a third of the saturated value. We have not studied what happens to a finite doping but a Fermi liquid regime on the hexagonal backbone at low doping looks likely. This would imply a conducting state, with the possibility of additional low-energy excitations in the form of defects of the hexagonal backbone.

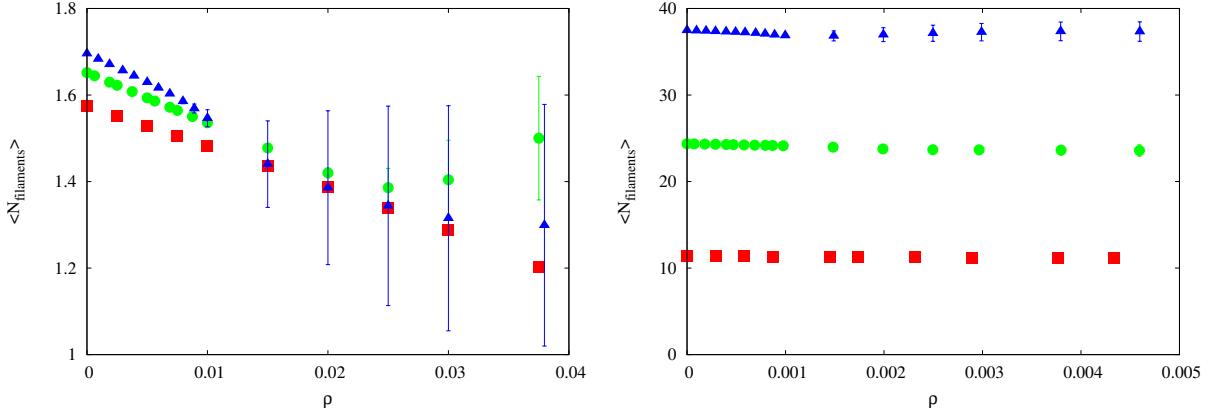


FIG. 12: Average number of filaments in 2D (*left*) and 3D (*right*), corresponding respectively to systems of sizes $L = 10, 20, 40$ and $L = 6, 12, 18$ (red squares, green dots, blue triangles) as a function of doping ρ .

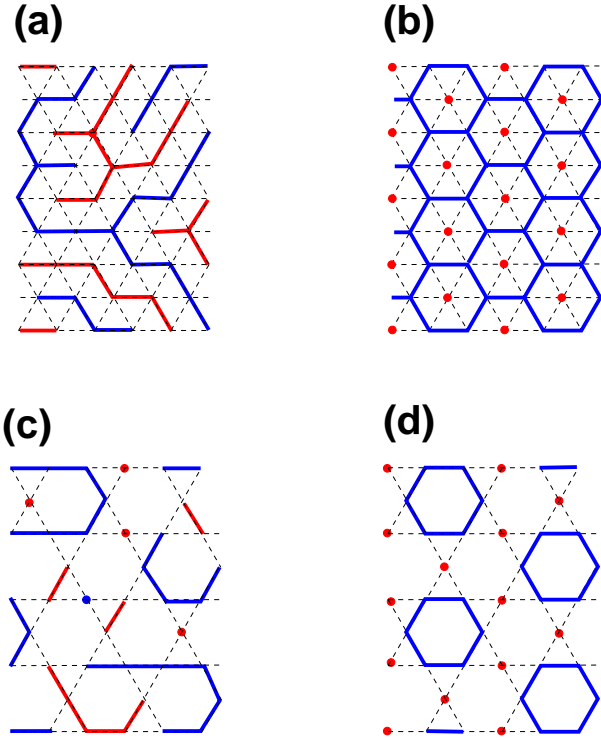


FIG. 13: Configurations on the triangular (top) and kagome (bottom) lattices respecting the antiferromagnetic frustrated constraints. (a) and (c) are random configurations, while (b) is the conducting hopping network minimizing the kinetic energy on the triangular lattice, and (d) is the insulating state maximizing the number of closed loops on kagome. Panel (d) also illustrates the arrangement of loops on kagome described in section IV C 2 in the two dimensional kagome layers which are part of the three dimensional pyrochlore lattice.

2. Kagome lattice

The situation on the kagome lattice is different still. Its magnetic ground state ensemble is very short-range correlated, unlike that of the triangular lattice, which has algebraic

correlations. With $z = 4$ for the kagome lattice, the hopping network can no longer branch but it can now have dangling links or isolated spins (coordination 1 and 0, respectively, as illustrated on Fig. 13.(c)).

The hopping networks thus consist of loop segments which need no longer close on themselves, minimally containing only one spin but not bounded above in the thermodynamic limit. The length distribution is, however, unlikely to contain long loops as this would only happen if each the magnetization of all the triangles along the loop segment has the same sign, at considerable cost in entropy.

Energetically, it is of course again most advantageous to have closed loops, as electrons on them gain hopping energy $-2|t|$. On the kagome lattice, such loops are readily constructed. The shortest ones are obtained by arranging spins to be aligned around a hexagon. Their densest packing is obtained for a state which breaks translational symmetry, tripling the unit cell and incorporating a $1/3$ magnetization as was the case of the triangular lattice above, see Fig. 13.(d). However, this state – which is the unique ground state at electron density $\rho = 1/9$ is now an insulating one – all hopping paths for the electrons close back on themselves after six steps.

Acknowledgments

We thank J. Chalker, C. Henley, and M. Viret for useful discussions.

Appendix A: Worm algorithm

a. Worm algorithm without electrons. In our definition, a loop is made of nearest neighbour spins pointing in the same direction: it is uniquely defined and possesses an up or down flavor. On the other hand, a worm consists of alternating up/down/up/down/... spins; through each vertex of the premedial lattice (square in 2D or diamond in 3D); see figure 14. The worm can randomly choose between two paths

energetically equivalent and eventually hit its initial position; the worm is then closed. Reversing all spins in the worm gives way to a new configuration in the Coulomb phase (figure 14). This method ensures both ergodicity and detailed balance to the algorithm in absence of electrons³⁶.

Worms, as defined above, are not self-avoiding, *i.e.* they can erase their own path. If a worm goes an odd number of times through the same site, the corresponding spin will be flipped during the Monte Carlo update; if it happens an even number of times, the spin shall not be flipped.




FIG. 14: On the left, the same loop configuration as in figure 1. The thick dashed green line represents a possible worm, as built by our algorithm; flipping the spins along this worm does not break the ice-rules (see on the right) and allows us to visit the configurational space of the Coulomb phase, ensuring ergodicity.

b. Decorrelation. Since we are dealing with Ising spins, two random configurations roughly differ by one half of the system sites on average. Hence two configurations separated by k worm updates can be considered decorrelated if close to one half of the system sites are different: we arbitrarily chose

45%. Previous work²⁰ has shown that a finite constant number of worm updates (~ 10) was enough to decorrelate the system in 3D, but an increasing number with system size L was necessary in 2D. This is related to Pólya's theorem stating that a random walk (similar to the worm here) in 3D is transient (finite probability never to come back to the origin), while it is recurrent in 2D (it always comes back)³⁷.

c. Worm algorithm with electrons – entropic regime. In the entropic regime, *i.e.* at low doping, there is at most one electron per loop. A correct sampling can then be done by using the above worm algorithm free of electrons, and rejecting all configurations with more electrons N_e than loops N_ℓ . However, the distribution of number of loops per configuration being Gaussian (see figure 3), the density of electrons acts as a cutoff and above a certain threshold of the order of $1/\langle \ell \rangle$, almost all configurations are rejected and it becomes impossible to get good statistics.

In order to take into account the influence of itinerant electrons in the entropy, we chose to weigh any loop configuration by the number of possible combinations to distribute N_e electrons in N_ℓ loops, namely the binomial coefficient $C_{N_\ell}^{N_e}$.

d. Worm algorithm with electrons – energetic regime. In the energy regime, *i.e.* at intermediate or high doping, there is more than one electron per loop. The above method based on the worm algorithm without electrons become inefficient and finite temperature Monte Carlo simulations are necessary.

Starting from a given configuration, a worm update is proposed. The loop histogram of the states before and after the proposed update are computed, and the corresponding electron energies are calculated after filling the energy levels with N_e electrons. The worm update is then accepted or not via a Metropolis argument.

¹ K. Ohgushi, S. Murakami, and N. Nagaosa, Physical Review B **62**, R6065 (2000).
² S. Onoda and N. Nagaosa, Physical Review Letters **90** (2003).
³ D. Ikoma, H. Tsuchiura, and J. Inoue, Physical Review B **68** (2003).
⁴ J. Haerter and B. Shastry, Physical Review Letters **95**, 087202 (2005).
⁵ Y. Shimomura, S. Miyahara, and N. Furukawa, Journal of the Physical Society of Japan **74**, 661 (2005).
⁶ I. Martin and C. D. Batista, Physical Review Letters **101**, 156402 (2008).
⁷ A. Kalitsov, B. Canals, and C. Lacroix, Journal of Physics: Conference Series **145**, 012020 (2009).
⁸ Y. Motome and N. Furukawa, Physical Review Letters **104** (2010).
⁹ S. Kumar and J. van den Brink, Physical Review Letters **105**, 216405 (2010).
¹⁰ M. Udagawa, H. Ishizuka, and Y. Motome, Physical Review Letters **104**, 226405 (2010).
¹¹ E. Tang, J.-W. Mei, and X.-G. Wen, Physical Review Letters **106**, 236802 (2011).
¹² D. A. Huse, W. Krauth, R. Moessner, and S. L. Sondhi, Physical Review Letters **91**, 167004 (2003).

¹³ G. Chen, J. Gukelberger, S. Trebst, F. Alet, and L. Balents, Physical Review B **80** (2009).
¹⁴ D. J. P. Morris, D. A. Tennant, S. A. Grigera, B. Klemke, C. Castelnovo, R. Moessner, C. Czternasty, M. Meissner, K. C. Rule, J. U. Hoffmann, et al., Science **326**, 411 (2009).
¹⁵ T. Fennell, P. P. Deen, A. R. Wildes, K. Schmalzl, D. Prabhakaran, A. T. Boothroyd, R. J. Aldus, D. F. McMorrow, and S. T. Bramwell, Science **326**, 415 (2009).
¹⁶ For a review of the Coulomb phase, see C. L. Henley, Annual Review of Condensed Matter Physics **1**, 179 (2010).
¹⁷ S. Powell, Physical Review B **84**, 094437 (2011).
¹⁸ J. Villain, Zeitschrift Fur Physik B-Condensed Matter **33**, 31 (1979).
¹⁹ S. T. Banks and S. T. Bramwell, arXiv:1107.5411 (2011).
²⁰ L. D. C. Jaubert, M. Haque, and R. Moessner, Physical Review Letters **107**, 177202 (2011).
²¹ A. Nahum, J. T. Chalker, P. Serna, M. Ortuno, and A. M. Somoza, Physical Review Letters **107**, 110601 (2011).
²² J. D. Bernal and R. H. Fowler, Journal of Chemical Physics **1**, 515 (1933).
²³ R. J. Baxter, Exactly solved models in statistical mechanics (Dover Publications, 2007).
²⁴ M. J. Harris, S. T. Bramwell, D. F. McMorrow, T. Zeiske, and K.

- W. Godfrey, *Physical Review Letters* **79**, 2554 (1997).
- ²⁵ R. Moessner, *Physical Review B* **57**, R5587 (1998).
- ²⁶ J. Cardy, *Annals of Physics* **318**, 81 (2005).
- ²⁷ J. L. Jacobsen and J. Vannimenus, *Journal of Physics A: Mathematical and General* **32**, 5455 (1999).
- ²⁸ M. Viret, F. Ott, J. Renard, H. Glatli, L. Pinsard-Gaudart, and A. Revcolevschi, *Physical Review Letters* **93** (2004).
- ²⁹ G.-W. Chern, *Physical Review Letters* **105**, 226403 (2010).
- ³⁰ R. Moessner, *Journal of Physics: Conference Series* **145**, 012001 (2009).
- ³¹ R. Moessner and S. L. Sondhi, *Physical Review B* **62**, 14122 (2000).
- ³² G. Murthy, D. Arovas, and A. Auerbach, *Physical Review B* **55**, 3104 (1997).
- ³³ D. Heidarian and K. Damle, *Physical Review Letters* **95**, 127206 (2005).
- ³⁴ S. Wessel and M. Troyer, *Physical Review Letters* **95**, 127205 (2005).
- ³⁵ R. Melko, A. Paramekanti, A. Burkov, A. Vishwanath, D. Sheng, and L. Balents, *Physical Review Letters* **95**, 127207 (2005).
- ³⁶ G. T. Barkema and M. E. J. Newman, *Physical Review E* **57**, 1155 (1998).
- ³⁷ B. D. Hughes, *Random walks and random environments* (Oxford University Press, 1995)