

# flatspin: A Large-Scale Artificial Spin Ice Simulator

Johannes H. Jensen,\* Anders Strømberg, Odd Rune Lykkebø,  
Arthur Penty, Magnus Sjölander, Erik Folven, and Gunnar Tufte  
*Norwegian University of Science and Technology, Trondheim, Norway*  
(Dated: December 18, 2020)

We present flatspin, a novel simulator for systems of interacting mesoscopic spins on a lattice, also known as artificial spin ice (ASI). Our magnetic switching criteria enables ASI dynamics to be captured in a dipole model. Through GPU acceleration, flatspin can simulate realistic dynamics of millions of magnets within practical time frames. We demonstrate flatspin’s versatility through the reproduction of a diverse set of established experimental results from the literature. In particular, magnetization details of “pinwheel” ASI during field-driven reversal have been reproduced, for the first time, by a dipole model. The simulation framework enables quick exploration and investigation of new ASI geometries and properties at unprecedented speeds. Finally, we use flatspin to explore new aspects of “square” ASI by introducing dilution defects and measuring the effect on the vertex population.

## I. INTRODUCTION

An artificial spin ice (ASI) is an ensemble of nanomagnets arranged on a lattice, coupled through magnetic dipole-dipole interactions. The vast variety of emergent collective behaviors found in these systems have generated considerable research interest over the last decade [12, 33]. Using modern nanofabrication techniques, emergent phenomena can be facilitated through direct control of the ASI geometry, e.g., collective ferromagnetic/antiferromagnetic ordering [34], Dirac strings [24], and phase transitions [2, 19]. ASIs offer a unique model system for exploring fundamental physics, since magnetic microscopy enables direct observation of their internal state. There is also a growing interest in ASIs as building blocks for novel devices [13, 14].

Micromagnetic simulations of ASI have been limited to a handful of nanomagnets due to excessive computational cost. Although physically accurate, such high fidelity simulations are unable to capture large-scale emergent phenomena, such as the size of magnetically ordered domains and long-range order. To simulate large ASI systems, an established approach is to sacrifice fidelity for speed by employing a dipole model, i.e., treating each nanomagnet as a single macro spin approximated by a point dipole [5]. Traditionally, Monte Carlo methods have been used in conjunction with the dipole approximation to search for low energy configurations [15, 39] or study statistical measures such as vertex populations [5]. However, Monte Carlo methods are inherently stochastic and better suited for ensemble statistics rather than detailed dynamics [3]. Monte Carlo methods will typically not provide the dynamical pathway taking the system from an initial configuration to the final low energy configuration.

flatspin is a simulator for large ASI systems based on a dipole approximation with the ability to capture re-

alistic dynamics, inspired by the work of Budrikis [5]. In contrast to Monte Carlo methods, the focus of flatspin is to simulate the dynamical pathways taken by ASI systems. The flexibility of flatspin enables quick exploration of ASI parameters and geometries. Through GPU acceleration, flatspin can capture realistic dynamics of millions of magnets within practical time frames.

In this paper, we present the motivation and design of flatspin. We demonstrate good agreement between flatspin and a variety of published experimental results. We show that flatspin can capture dynamic behaviors observed experimentally, which have previously eluded modeling [20]. Finally, we explore new aspects of square ASI by removing individual elements of the lattice, and measuring the effects on the type population under an applied field.

## II. THE FLATSPIN MAGNETIC MODEL

In this section, we describe the dipole model and the underlying physical assumptions of flatspin. The model is designed to simulate the ensemble state-by-state evolution, i.e., *dynamics*, of two-dimensional ASI. In short, magnets are modeled as point dipoles (section II A), and each dipole can be affected by three types of external influence: magnetic dipole-dipole coupling (section II C), an applied external magnetic field (section II D), and thermal fluctuations (section II E). The switching of spins is determined using a generalized Stoner-Wohlfarth model (section II F). Imperfections in the ASI are introduced as different coercive fields, set per spin (section II G). Dynamics are modeled using a deterministic single spin flip strategy (section II H).

### A. Magnets as dipoles

ASI systems are physically realized as elongated islands of a ferromagnetic material, arranged on a two-dimensional lattice. The magnets are made small enough

---

\* johannes.jensen@ntnu.no




FIG. 1. The representation of nanomagnets as spins  $s_i$  and associated quantities: angle  $\theta_i$  and distance to neighbor  $j$ ,  $\mathbf{r}_{ij}$ . Note that the magnetization of spin  $i$  is given by its spin,  $s_i$ , and orientation,  $\theta_i$ .

to exhibit a single ferromagnetic domain, i.e., coherent magnetization throughout the magnet. The single domain is stable as the energy cost associated with domain walls exceeds the cost associated with the demagnetization energy [9, 17]. Since a magnet has coherent magnetization, it can be approximated by a single mesoscopic spin and the magnetic state can be represented by a single vector  $\mathbf{m}$ .

The magnets will exhibit an in-plane shape anisotropy due to their small thickness, as well as a magnetization direction defined by their elongated shape. Hence individual magnets can be approximated by classical macro spins with a twofold degenerate ground state defined by the elongated shape of the individual elements. Due to the two degenerate ground state configurations, we approximate each magnet as a magnetic dipole with *binary* magnetization, i.e., a macro spin,  $s_i \in \{-1, +1\}$ .

As illustrated in Fig. 1, each magnetic dipole is modelled with a position  $\mathbf{r}_i$  and rotation  $\theta_i$ , which together define the ASI geometry. Furthermore, each magnet is assigned a coercive field,  $h_c^{(i)}$ , describing its resistance to switching (see section IIF). Using reduced units, the magnetization vector of a single magnet can be expressed as

$$\mathbf{m}_i = s_i \hat{\mathbf{m}}_i \quad (1)$$

where  $\hat{\mathbf{m}}_i$  is the unit vector along  $\mathbf{m}_i$ .

## B. Magnetic Fields and Temperature

External fields and temperature are modeled as a combination of effective magnetic fields. The total field,  $\mathbf{h}_i$ , affecting each magnet  $i$  is the sum of three components:

$$\mathbf{h}_i = \mathbf{h}_{\text{dip}}^{(i)} + \mathbf{h}_{\text{ext}}^{(i)} + \mathbf{h}_{\text{th}}^{(i)}, \quad (2)$$

where  $\mathbf{h}_{\text{dip}}^{(i)}$  is the local magnetic field from neighboring magnets (magnetic dipole-dipole interactions),  $\mathbf{h}_{\text{ext}}^{(i)}$  is a global or local external field, and  $\mathbf{h}_{\text{th}}^{(i)}$  is a random magnetic field representing thermal fluctuations in each magnetic element. The contributions from each field is described in the following sections.

## C. Magnetic dipole-dipole interactions

The individual magnets, or spins, are coupled solely through dipole-dipole interactions. Each spin,  $i$ , is subject to a magnetic field from all neighboring spins,  $j \neq i$ , given by

$$\mathbf{h}_{\text{dip}}^{(i)} = \alpha \sum_{j \neq i} \frac{3\mathbf{r}_{ij}(\mathbf{m}_j \cdot \mathbf{r}_{ij})}{|\mathbf{r}_{ij}|^5} - \frac{\mathbf{m}_j}{|\mathbf{r}_{ij}|^3}, \quad (3)$$

where  $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$  is the distance vector from spin  $i$  to  $j$ , and  $\alpha$  scales the dipolar coupling strength between spins. The coupling strength  $\alpha$  is given by  $\alpha = \frac{M}{4\pi a^3}$ , where  $a$  is the lattice spacing, and  $M$  is the absolute magnetic moment of a single magnet. The distance  $\mathbf{r}_{ij}$  is given in reduced units of the lattice spacing.

The dipole field present at each spin's location is calculated by summing the dipole field contributions from spins in its neighborhood. The size of the neighborhood is user-configurable and defined in units of the lattice spacing. The required neighborhood distance varies, subject to the system of study. Care must be taken to include enough spins in the neighborhood such that the observed behavior converges, especially when considering systems exhibiting long-range effects. In some geometries, such as square ASI, short range interactions dominate the contributions to  $\mathbf{h}_{\text{dip}}$  [6, 38], in which case the neighborhood size can be relatively small, for a benefit of increased efficiency. For geometries where long range interactions are significant, a larger neighborhood is required, e.g., pinwheel ASI [21].

## D. External field

Applying an external magnetic field is the primary mechanism for altering the state of an ASI in a controlled manner. The external field can either be set locally on a per-spin basis,  $\mathbf{h}_{\text{ext}}^{(i)}$ , globally for the entire system,  $\mathbf{h}_{\text{ext}}$ , or as a spatial vector field,  $\mathbf{h}_{\text{ext}}(\mathbf{r})$ .

Time-dependent external fields are supported, i.e.,  $\mathbf{h}_{\text{ext}}$  is a discrete time series of either local, global or spatial fields. A variety of time-dependent external fields are provided, including sinusoidal, sawtooth and rotational fields. More complex field-protocols can be generated, e.g., for annealing purposes or probing dynamic response.

### E. Thermal field

In a physical ensemble of particles, there will be a random distribution of thermal energy, on the scale of  $k_B T$ , where  $T$  is temperature and  $k_B$  is the Boltzmann constant. This thermal energy will cause random switching events in an ensemble of bistable magnetic particles, at a characteristic rate given by the Néel-Brown equation [4, 25],

$$f = f_0 \exp\left(-\frac{\Delta E}{k_B T}\right), \quad (4)$$

where  $f_0$  is the attempt frequency and  $\Delta E$  the particle's energy barrier for switching.

During a finite time interval,  $\Delta t = 1/f$ , defining the experimental time scale, the expected amount of energy provided to a particle by thermal fluctuations can be expressed as

$$\Delta E = k_B T \ln(f_0 \Delta t). \quad (5)$$

In flatspin we represent the fluctuations in thermal energy,  $E_{\text{th}}$ , by a corresponding magnetic field. The magnitude of this field is such that its Zeeman energy is equal to the thermal energy contribution for a single particle,  $H_{\text{th}} = E_{\text{th}}/M$ , where  $M$  is the magnetic moment.

We independently draw the components of this field from a Normal distribution, with mean zero and standard deviation  $\sigma_{\text{th}}$ . For two-dimensional vectors, the magnitude will then follow a Rayleigh distribution with a mean  $\mu = \sigma_{\text{th}} \sqrt{\frac{\pi}{2}}$ .

By setting

$$\sigma_{\text{th}} = \sqrt{\frac{2}{\pi}} \frac{k_B T}{M_S V} \ln(f_0 \Delta t), \quad (6)$$

where  $M_S$  the saturation magnetization,  $V$  the volume of the magnetic islands, the expected magnitude of the thermal field fluctuation is such that its Zeeman energy corresponds to that of  $\Delta E$  in Eq. (5). In flatspin, the user can directly set  $\sigma_{\text{th}}$ , or provide the parameters in Eq. (6). Thus, thermal fluctuations are modeled as an additional local field,  $\mathbf{h}_{\text{th}}^{(i)}$ , applied to each magnet individually.

It is important to note, flatspin does not account for the temperature dependence of the parameters in Eq. (5). If these parameters are expected to vary significantly in the temperature range of interest, e.g.,  $M_S$ , this has to be explicitly accounted for by the user.

### F. Switching

Magnetization reversal, or *switching*, may take place when a magnet is subjected to a magnetic field or high temperature. If the field is sufficiently strong (stronger than some critical field) and directed so that the projection onto  $\mathbf{m}_i$  is in the opposite direction to  $\mathbf{m}_i$ , the magnetization  $\mathbf{m}_i$  will switch direction.

The critical field strength is referred to as the coercive field  $h_c$ . For elongated magnets,  $h_c$  depends on the angle between the applied field  $\mathbf{h}_i$  and  $\mathbf{m}_i$ . As illustrated in Fig. 2a, the *easy axis*, where the magnetization favors alignment, lies along the long axis of the magnet, whereas the *hard axis* is perpendicular to the long axis. The external field can be decomposed into two components,  $\mathbf{h}_{\parallel}$  and  $\mathbf{h}_{\perp}$ , corresponding to the field component parallel and perpendicular to the easy axis, respectively. We denote the coercive field strength along the hard axis as  $h_k$ .

A *switching astroid* is a polar plot of  $h_c$  at different angles, with  $h_{\perp}$  on the horizontal axis and  $h_{\parallel}$  on the vertical axis. It is a description of  $\mathbf{h}_c$ , or  $h_c(\phi)$ . For any applied field  $\mathbf{h}_i$  that is outside the switching astroid, the magnet will switch (given that the projection of  $\mathbf{h}_i$  onto  $\mathbf{m}_i$  is oppositely aligned with respect to  $\mathbf{m}_i$ ).

Fig. 2c shows the normalized switching astroid for an elliptical magnet (Fig. 2a) as obtained from micromagnetic simulations (red dots). Notice how  $h_c$  is the smallest at a 45° angle, indicating that a field directed at 45° to a magnet's principal axes will require the least field strength in order to switch its magnetization.

The Stoner-Wohlfarth (SW) model captures the angle dependent switching characteristic of single-domain elliptical magnets [36]. The characteristic SW astroid is shown in Fig. 2c (blue line) and is described by the equation

$$\left(\frac{h_{\parallel}}{h_k}\right)^{2/3} + \left(\frac{h_{\perp}}{h_k}\right)^{2/3} = 1. \quad (7)$$

In the SW model, switching may occur when the left hand side of Eq. (7) is greater than one.

The astroid obtained from micromagnetic simulations and the SW astroid (Fig. 2c) are nearly identical, indicating that the SW model is a simple and valid description of switching in elliptical nanomagnets.

However, the SW model is only accurate for elliptical magnets, other magnet shapes typically have quite different switching characteristics. Fig. 2d shows the switching astroid for rectangular stadium-shaped magnets (red dots), which is the shape commonly used in most fabricated ASIs (Fig. 2b). Notice how the astroid is asymmetric: rectangular magnets switch more easily along the easy axis than the hard axis.

To capture the asymmetric switching characteristics of non-elliptical magnets, we have generalized the SW switching model to allow an asymmetry between easy and hard axes. Additionally, the model has been extended to support tuning of the curvature of the extrema. In the generalized model, the switching threshold is given by

$$\left(\frac{h_{\parallel}}{bh_k}\right)^{2/\gamma} + \left(\frac{h_{\perp}}{ch_k}\right)^{2/\beta} = 1, \quad (8)$$

where  $b$ ,  $c$ ,  $\beta$  and  $\gamma$  are parameters which adjust the shape of the astroid:  $b$  and  $c$  define the height and width,




FIG. 2. Top: Schematic showing hard and easy axes of (a) an elliptical magnet and (b) a rectangular stadium-shaped magnet, as well as the total field acting on the magnet,  $\mathbf{h}_i$ , with its parallel and perpendicular components,  $\mathbf{h}_\parallel$  and  $\mathbf{h}_\perp$ , respectively. Bottom: Switching astroid for (c) an elliptical magnet and (d) a rectangular stadium-shaped magnet. Red dots show the coercive field obtained from micromagnetic simulations. The blue line in (c) shows the Stoner-Wohlfarth astroid. The blue line in (d) shows the generalized Stoner-Wohlfarth astroid with parameters  $b = 0.42$ ,  $c = 1$ ,  $\beta = 1.7$ , and  $\gamma = 3.4$  in Eq. (8). The astroids have been normalized with respect to  $h_k$ .

respectively, while  $\beta$  and  $\gamma$  adjust the curvature of the astroid at the easy and hard axis, respectively. Introducing these new parameters allows for tuning of the switching astroid to fit with the shape of nanomagnets used in ASIs. With  $b = c = 1$  and  $\beta = \gamma = 3$ , Eq. (8) reduces to Eq. (7), i.e., the classical Stoner-Wohlfarth astroid is obtained (valid for elliptical magnets).

By tuning the parameters of the generalized SW model, we can obtain the asymmetric switching astroid shown in Fig. 2d (blue line). The astroid is in good agreement with results obtained from micromagnetic simulations (red dots).

In flatspin, the generalized SW model is used as the switching criteria, i.e., a spin may flip if the left hand side of Eq. (8) is greater than one. Additionally, the projection of  $\mathbf{h}_i$  onto  $\mathbf{m}_i$  must be in the opposite direction of  $\mathbf{m}_i$ :

$$\mathbf{h}_i \cdot \mathbf{m}_i < 0. \quad (9)$$

### G. Imperfections and disorder

Due to manufacturing imperfections there will always be a degree of variation in the shape and edge roughness of nanomagnets. This variation can be thought of

as a disorder in the magnets' inherent properties. Rough edges and a slightly distorted geometry can affect the magnets' switching mechanisms, with defects pinning magnetization and altering the coercive field for each magnet.

In flatspin we model this variation as disorder in the coercive fields. The coercive field is defined individually for each magnet, and a distribution of values can be used to introduce variation. A user-defined parameter,  $k_{\text{disorder}}$ , defines the distribution of coercive fields, i.e.,  $h_k^{(i)}$  is sampled from a normal distribution  $\mathcal{N}(h_k, \sigma)$ , where  $\sigma = k_{\text{disorder}} \cdot h_k$  (while ensuring  $h_k^{(i)}$  is always positive).

### H. Dynamics

flatspin employs deterministic single spin flip dynamics. At each simulation step, we calculate the total magnetic field,  $\mathbf{h}_i$ , acting on each spin. Next, we determine which spins *may* flip according to the switching criteria Eqs. (8) and (9). Finally we flip the spin where  $\mathbf{h}_i$  is the furthest outside its switching astroid, i.e., where the left hand side of Eq. (8) is greatest. Ties are broken in a deterministic, arbitrary manner, although with non-zero




FIG. 3. flatspin includes the most common ASI geometries: (a) Square (closed edges), (b) Square (open edges), (c) Kagome, (d) Pinwheel “diamond”, (e) Pinwheel “lucky-knot”, and (f) Ising.

disorder such occurrences are rare. The dipolar fields are recalculated after every spin flip and the above process is repeated until there are no more flippable spins.

This relaxation process is performed with constant external and thermal fields. To advance the simulation, the fields are updated and relaxation is performed again. Hence a simulation run consists of a sequence of field updates and relaxation processes.

The dynamical process makes three main assumptions:

1. The external field is quasi-static compared to the time scale of magnet switching.
2. Magnet switching is sequential.
3. The magnet experiencing the highest effective field compared to its switching threshold is the first to switch.

*Assumption 1* means the model holds for low frequency external fields, i.e., where switching settles under a static field. The switching mechanics of nanomagnets are typically in the sub nanosecond range [11, 16], and experimental setups often employ external magnetic fields which can be considered static at this time scale. At high applied field frequencies, more complex physical phenomena such as spin waves will have a non-negligible effect on switching dynamics, which are not modeled in flatspin.

*Assumption 2* holds if the coercive fields  $h_c^{(i)}$ , and total field  $\mathbf{h}_i$ , of the magnets are sufficiently non-uniform, so that there will always be a single magnet which will flip first. It is assumed to be unlikely that two magnets will have the same  $h_c^{(i)}$  and  $\mathbf{h}_i$  simultaneously. However, in those rare cases where two magnets are equally close to switching, overlapping switching events may occur in a physical system.

*Assumption 3* relies on the fact that all changes in the magnetic fields are effectively continuous, and the change is unidirectional within a simulated time step, i.e., a *quasi-static field*. Since evaluation happens in discrete time, there will be cases where several magnets are above their corresponding switching thresholds simultaneously. In those cases, the magnet furthest above its switching threshold will have been the first to have crossed the threshold under a quasi-static field. Furthermore, if the angle of the external field is constant, the switching order is invariant to the time resolution of the external field.

## I. Geometries

The particular spatial arrangement of the magnets is referred to as the *geometry*. A wide range of ASI ge-




FIG. 4. Overview of the flatspin architecture, with arrows indicating data flow.

ometries have been proposed in the literature. Fig. 3 depicts the geometries included in flatspin, which are the most commonly used ASI geometries: square [38], kagome [28, 35], pinwheel [20, 21] and ising [30]. Note that when we refer to “pinwheel ASI” in this work, we are referring explicitly to the  $45^\circ$  variant.

Geometries are often decomposed into two or more “sublattices”, where the magnets within one sublattice are all aligned (have the same rotation). In Fig. 3, the sublattice a magnet belongs to is indicated by its color. As can be seen, both square and pinwheel ASIs have two perpendicular sublattices, whereas kagome has three sublattices.

flatspin can be used to model any two-dimensional ASI comprised of identical elements. New geometries can easily be added by extending the model with a new set of positions  $\mathbf{r}_i$  and rotations  $\theta_i$ .

### III. SIMULATION FRAMEWORK

In addition to a magnetic model, flatspin provides a flexible framework for running simulations, storing results, and performing analysis.

Fig. 4 illustrates the overall architecture of flatspin. The *ASI model* has been described in detail in section II. Conceptually, the ASI model describes the physical system under study. The rest of the components are tools used by the experimenter to interact with the ASI and observe the results. In this section we briefly describe each of these components.

The *input encoder* translates a set of input values to a series of external fields. Encoders provide a flexible way to define field protocols, and have been designed with neuromorphic computing in mind. A range of encoders are included, e.g., sinusoidal, sawtooth and rotational fields.

The responsibility of the *runner* component is to perturb the ASI model according to the field protocol, and save the results. The model, which is fully parametric, receives parameters from the runner, enabling automated parameter sweeps. In addition, there is support for distributed running of simulations on a compute cluster.

Results are stored in a well-defined *dataset* format which makes the analysis of a large numbers of simulations straightforward. A suite of *analysis tools* are included, e.g., for plotting results, visualizing ensemble dynamics and analysis of vertex populations.

flatspin is written in Python and utilizes OpenCL to accelerate calculations on the GPU. OpenCL is supported by most GPU vendors, hence flatspin can run accelerated on a wide variety of platforms. The simulator may also run entirely on CPUs in case GPUs are not available, albeit at a reduced speed.

flatspin is open-source software and released under a GNU GPL license. For more information see the website [8] and User Manual [7].

## IV. VALIDATION OF FLATSPIN

To evaluate the ASI model, flatspin simulations were compared to established experimental results from the literature, as well as micromagnetic simulations. In particular we investigate phenomena such as Dirac strings in kagome ASI, temperature-dependent superparamagnetic Langevin dynamics, the size of crystallite domains in square ASI, and superferromagnetism in pinwheel ASI. Finally, we compare the switching order from flatspin simulations with that of micromagnetic simulations, and investigate the effect of varying lattice spacings.

### A. Dirac strings in kagome ASI

To assess the ability of flatspin to reproduce fine-scale patterns, we consider the emergence of Dirac strings in a kagome ASI (Fig. 3c). Applying a reversal field to a polarized kagome ASI results in the formation of monopole-antimonopole pairs [22]. These pairs are joined by a “string” of nanomagnets which have flipped due to the reversal field. As the strength of the reversal field increases, the strings elongate until they fill the array.

We closely follow the methodology set out in an experimental study of Dirac strings in kagome ASI [22], in which a room temperature kagome ASI undergoes magnetization reversal. We start with an array of 2638 magnets ( $29 \times 29$  hexagons) polarized to the left and apply a reversal field  $\mathbf{H}$  to the right with a slight, downward offset of  $3.6^\circ$ . This offset breaks the symmetry, such that one of the sublattices is now least aligned with the field, resulting in an increased coercive field on this “unfavored” sublattice.

Micromagnetic simulations of magnets of size  $470 \text{ nm} \times 160 \text{ nm} \times 20 \text{ nm}$  yield the following estimation



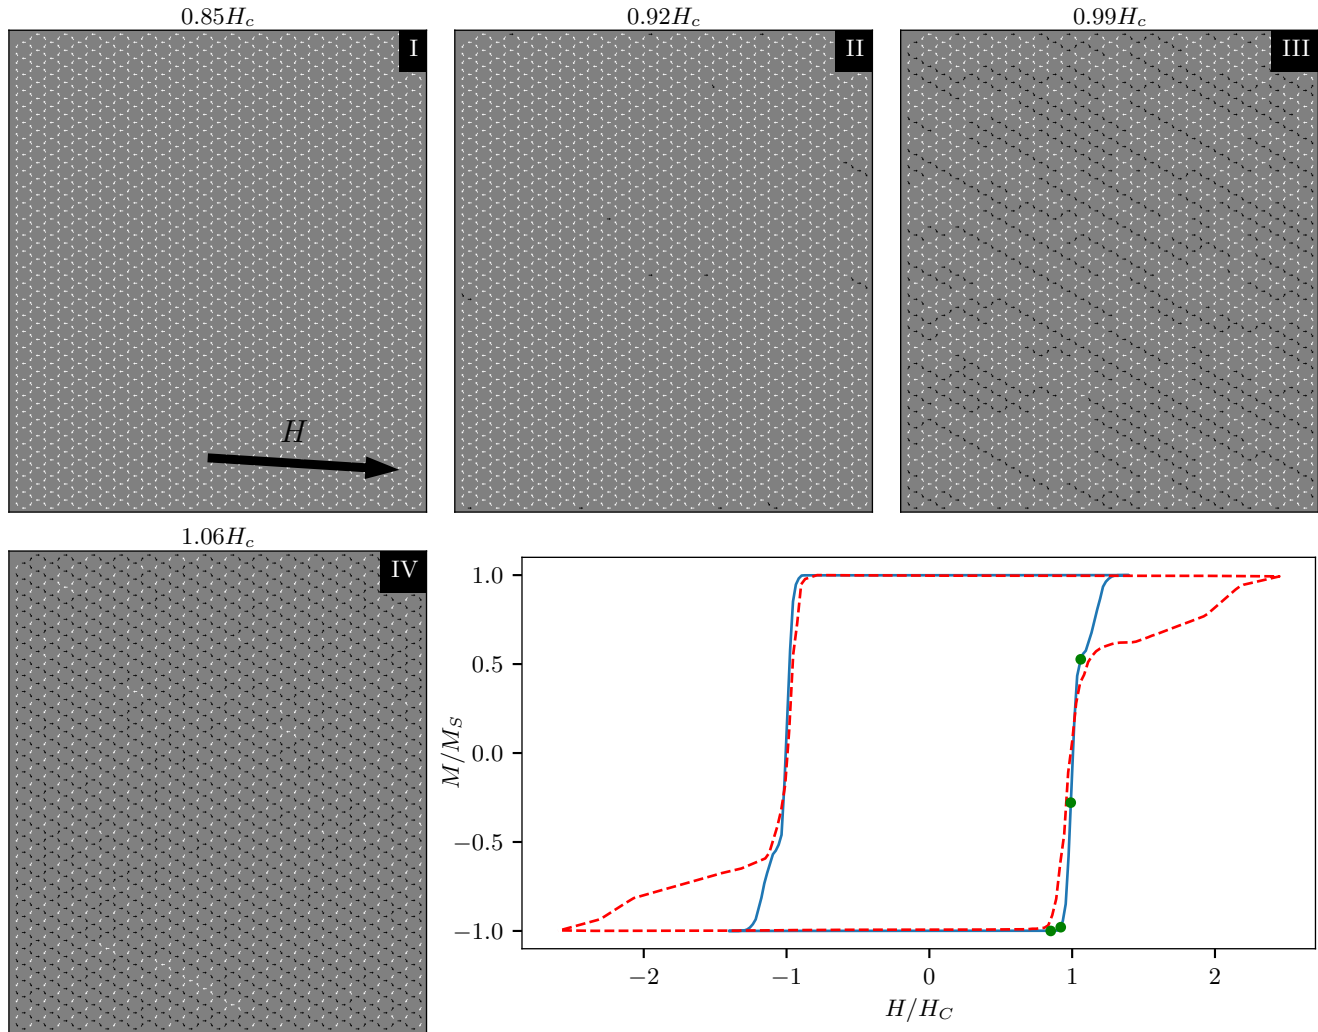


FIG. 5. Left: snapshots of the evolution of a kagome ASI at selected field values. Right: Comparison of the hysteresis curve of the simulated ensemble (blue line) against a sketch of the hysteresis curve from the experimental results [22] (red dashed line). Green dots indicate the points at which the snapshots are sampled from.

of flatspin parameters:  $\alpha = 0.00103$ ,  $h_k = 0.216$ ,  $\beta = 2.5$ ,  $\gamma = 3$ ,  $b = 0.212$ ,  $c = 1$ . The ensemble was simulated at constant room temperature (300 K) with timestep length  $\Delta t = 60$  s. The disorder parameter  $k_{\text{disorder}} = 0.07$  was determined empirically through qualitative comparison with the experimental results [22].

The time evolution snapshots of Fig. 5 demonstrate a strong, qualitative similarity to the results of Mengotti et al. [22]. We see Dirac strings developing with a preference to lie along the two sublattices most aligned with the field angle. Furthermore, in the final image, we see the vast majority of unflipped magnets (excluding the edges) are on the unfavored sublattice, in accordance with both experimental and simulated results from the literature.

Also in Fig. 5, we see the hysteresis of the simulated ensemble (solid line) is similar to that of Mengotti et al. [22] (dashed line) in some sections, but differs near the extrema. The hysteresis can be understood in two phases.

The first phase, at roughly  $M/M_S \in [-0.6, 0.6]$ , is dominated by the lengthening of the Dirac strings, with almost no activity occurring on the unfavored sublattice. At  $M/M_S < -0.6$  and  $M/M_S > 0.6$ , the ensemble enters a second phase in which the Dirac strings have fully covered the array, and change in magnetization is dominated by switching on the unfavored sublattice. Clearly we see good agreement, within phase one, between our simulated hysteresis and the experimental results. Furthermore, there is a clear phase transition (characterized by a sharp decrease in gradient) in our hysteresis very close to the transition in the experimental hysteresis. Notably however, although the phase transitions occur at similar field values, the change in gradient is less pronounced in our simulated hysteresis. This disparity indicates that, in the second phase, the magnets on the unfavored sublattice flip more easily in our simulation than in the experimental data.

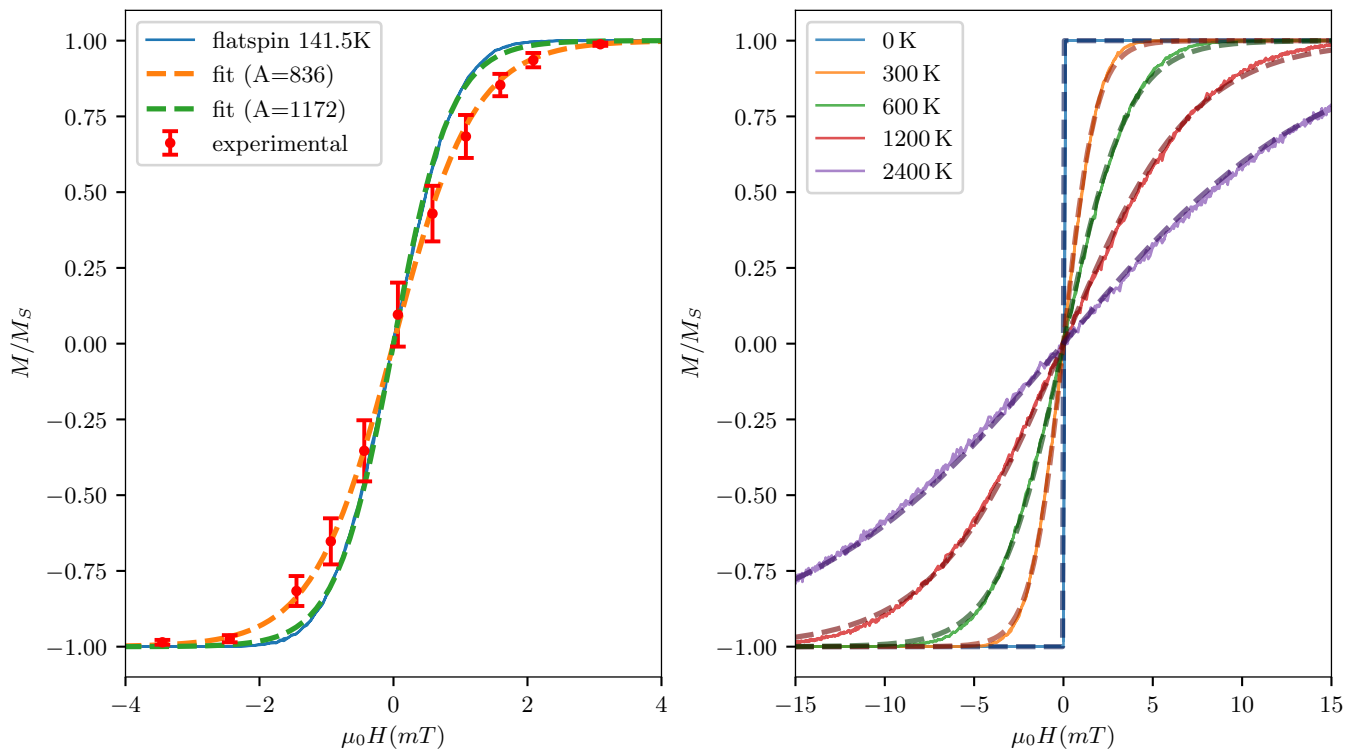


FIG. 6. Left: flatspin simulations corresponding to the system of non-interacting SrRuO<sub>3</sub> nanomagnets measured by Sinwani et al. [31], along with their experimentally obtained data of the magnetization density. Dashed lines show fits to Eq. (10) using the fitting parameter  $A = 836 \text{ T}^{-1}$  from the original work and  $A = 1172 \text{ T}^{-1}$ , the best fit for the flatspin data. Right: flatspin simulations (solid lines) of the corresponding system, obtained at different temperatures together with Eq. (10) (dashed lines), using the fitting parameter,  $A = 1172 \text{ T}^{-1}$ , found for  $T = 114.5 \text{ K}$ , but scaled inversely by temperature ( $A \propto T^{-1}$ ). All flatspin simulations are averaged over one hundred full hysteresis loops.

The accuracy of the point dipole approximation is known to suffer when considering kagome ASI. Specifically, it has been shown to underestimate the coupling coefficient of the nearest neighbors by approximately a factor of 5 [29]. Despite this, we observe flatspin accurately reproduces snapshots of the time evolution behavior, while also capturing salient features of the ensemble hysteresis curve.

### B. Superparamagnetic Langevin dynamics

One method to examine the correctness of flatspin’s implementation of temperature, is to simulate a system simple enough that the expected behavior can be known analytically. To this end, we investigate the effect of temperature on a system of non-interacting spins at thermal equilibrium, and compare it to experimental results.

Sinwani et al. [31] have investigated superparamagnetic switching in uncoupled SrRuO<sub>3</sub> nanomagnets, using anomalous Hall effect measurements on single magnets. To verify the implementation of temperature in our simulator, we compare flatspin’s results with their experimental data.

For non-interacting spins the coupling coefficient is

zero, meaning each spin is independent. In a superparamagnetic system, the switching threshold is small compared to the thermal energy, so its state is determined only by the external field and thermal noise. In flatspin such a system is modeled by setting  $\alpha = 0$  and  $h_k = 0$ , and using the ‘Ising’ geometry shown in Fig. 3f.

In Fig. 6, we compare the experimental data against flatspin simulations of the corresponding system. The parameters used for the simulations matching the experiments by Sinwani et al. [31], were  $V = 2.093 \times 10^{-22} \text{ m}^3$ ,  $M_S = 213 \times 10^3 \text{ A m}^{-1}$ , and  $\Delta t = 1 \text{ s}$  for the macrospin volume, saturation magnetization, and measurement time, respectively. For the purpose of comparison, the bias of 1.06 mT, originating from remanent magnetization in their superconducting coil and ambient fields, was removed from the experimental data. We find good agreement between the experimental data (Fig. 6). In the simulations the response to the external field is slightly stronger. This could be explained by magnetic pinning in the physical system, or a small offset in temperature in either flatspin or in the experiment.

In the original paper by Sinwani et al. [31], the authors provide a fit to an analytical equation,

$$M/M_s = \tanh(A(\mu_0 H + B)), \quad (10)$$



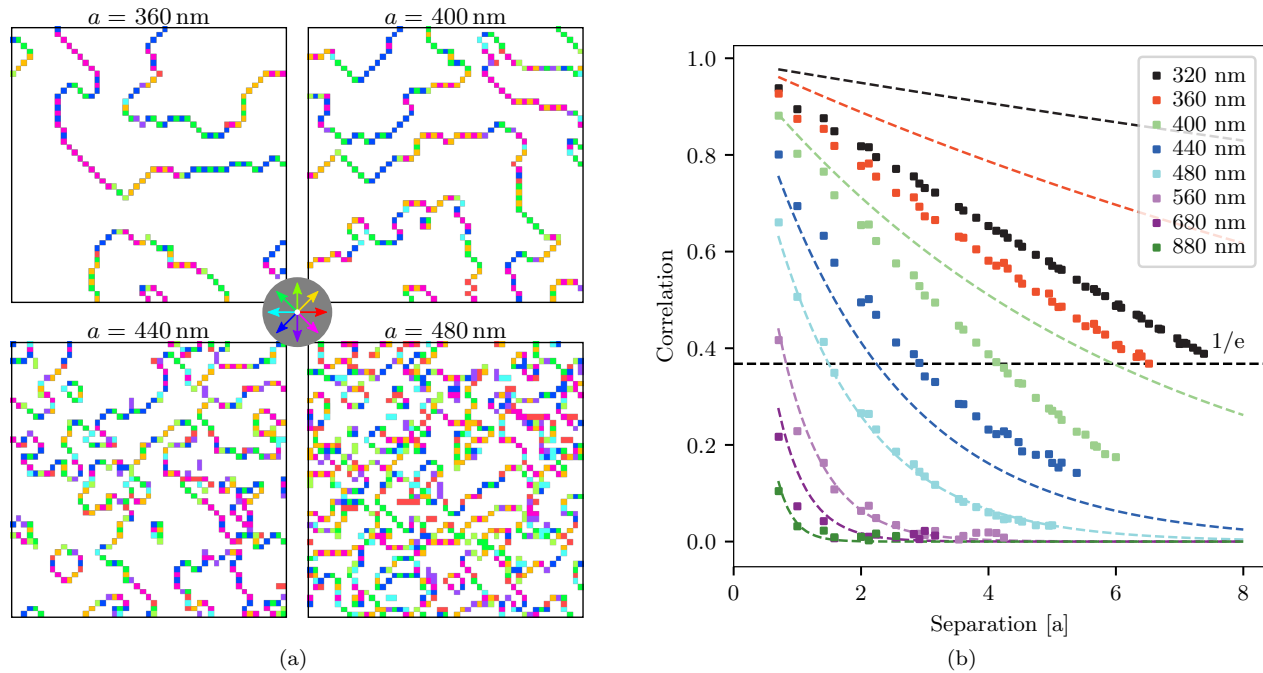


FIG. 7. (a) Maps showing the net magnetization of the ASI vertices, for an annealed  $50 \times 50$  square ASI with the given lattice constant  $a$ . The white regions have zero net magnetization and thus correspond to a coherent domain of Type I vertices. Colored regions have a non-zero net magnetization, direction indicated by the color wheel, and correspond to Type II or III vertices. (b) The absolute value of spin-spin correlation at a given separation for square ASIs of different lattice constants,  $a$ . Dashed curves show exponential best fits for data from the original paper [39]. The significant deviation for the results of the 320 nm lattice constant system is discussed in the text. Also indicated is a  $1/e$  threshold of correlation.

where  $A$  is a scaling factor and  $B$  a remanent bias field magnitude. The authors found the best fit at values  $A = 836 \text{ T}^{-1}$  and  $B = 1.06 \text{ mT}$ . For the simulated data the best fit values were  $A = 1172 \text{ T}^{-1}$ , and since there are no bias fields,  $B = 0 \text{ T}$ .

To check if temperature scales correctly in flatspin, we assume that  $A$  scales with  $1/T$ , and compare simulations at different temperature against Eq. (10). We find good agreement between the analytical expression and flatspin results, indicating that flatspin scales temperature in a physical manner. Although the presented results were quite close to experimentally obtained values, we note that the switching mechanisms in nanomagnets can differ significantly between systems. Therefore, a calibration of  $\sigma_{\text{th}}$  for the system at hand may be necessary. However, as flatspin has been shown to scale temperature in a physically realistic manner, see Fig. 6, calibration at one temperature should suffice.

### C. Domain size in square ASI

In order to demonstrate simulation of large-scale behavior, we have reproduced the emergence of large domains of magnetic order in square ASI, similar to experimental results of Zhang et al. [39]. One of the main advantages of flatspin over typical alternatives is the scal-

ability and high throughput of large systems with many magnets. Some emergent ASI phenomena require large systems in order to be fully quantified and studied with high fidelity, such as the domain size of magnetic charge crystallites. For ASIs with strongly coupled magnets, typical domain sizes can become too large for direct experimental observation. Thus, an accurate estimate of the domain size for ASIs with a small lattice constant is, in part, limited by the number of directly observable magnets.

For a given range of lattice constants covering both strongly coupled ASIs and weakly coupled ASIs, a corresponding range of large to small magnetic order coherence lengths is expected. In this study, we consider square ASI (closed edges, Fig. 3a) with different lattice constants,  $a$ , ranging from 320 nm to 880 nm.

$50 \times 50$  square ASIs were annealed in flatspin using a thermal protocol of exponentially decreasing  $\sigma_{\text{th}}$ . A switching astroid for  $220 \text{ nm} \times 80 \text{ nm} \times 25 \text{ nm}$  was obtained through micromagnetic simulations, described by generalized astroid parameters  $b = 0.38$ ,  $c = 1.0$ ,  $\beta = 1.3$ , and  $\gamma = 3.6$ . Additionally,  $h_k = 0.20$ ,  $k_{\text{disorder}} = 0.05$ , and a neighbor distance of 10 lattice constants was used. The thermal protocol was chosen such that the total dipole interaction energy was not significantly reduced by increasing simulation steps.

In the annealed state, the spin-spin correlation as a

function of their lateral separation was calculated across the ensembles. Analysis of the average correlation of annealed states provides insight about the typical coherence length of magnetic order, i.e., magnetic charge crystallite size, or domain size. Here, the correlation of two spins is defined as +1 (-1) if their dipole interaction is minimized (maximized). Averaging correlation across distinct types of spin pairs, in the annealed ASI, gives a measure of how coherent the ASI is at that particular neighbor separation. How quickly the average correlation decreases as a function of separation can be used to estimate the characteristic domain size. In particular, it can be argued that the separation where the correlation falls below  $1/e$  is the characteristic domain radius [26, 39].

Typical domain structures and correlation results can be seen in Fig. 7. The domains shown in Fig. 7a and the correlation curves in Fig. 7b are, for the most part, in good agreement with experimental results [39]. A qualitative comparison of the domain sizes and structures in Fig. 7 shows that the domains tend to be larger, with smoother domain boundaries, for smaller  $a$ . The analysis of coherence as a function of separation also shows identical trends and similar values, where an increase in  $a$  leads to low correlation, even between nearest neighbors.

For cases where  $a < 400$  nm, domains do not significantly increase in size when the lattice constant is further reduced. This discrepancy with the experimental results is not completely unexpected: the point-dipole approximation is known to underestimate nearest-neighbor interaction for magnets placed close together [29]. In addition, a stronger interaction between spins would cause each spin flip to contribute a greater change in the total dipole energy. This makes a gradual descent towards the ground state by random spin flips (the thermal fluctuations as modeled by flatspin) harder to achieve. These issues may be addressed by increasing the coupling parameter  $\alpha$  for nearest neighbor spins, and by a longer and slower annealing protocol. A longer and slower annealing protocol will inevitably come at the cost of longer computation times. In a future version of flatspin, it might be beneficial to include other thermal effects such as a declining saturation magnetization of the constituent material.

These results show that flatspin provides sufficient flexibility, fidelity and performance required to reproduce experimentally observed large-scale emergent behavior in ASIs.

#### D. Superferromagnetism in pinwheel ASI

In this section, we use flatspin to reproduce the dynamic behavior of pinwheel ASI, which had yet to be demonstrated with a dipole model [20]. We find that our switching criteria plays a key role in replicating magnetization details during the field-driven array reversal.

Pinwheel ASI is obtained by rotating each island in square ASI some angle about its center. A rotation of 45

degrees results in a transition from antiferromagnetic to ferromagnetic order [21]. The dynamics of pinwheel ASI in many ways resemble continuous ferromagnetic thin films, with mesoscopic domain growth originating from nucleation sites, followed by coherent domain propagation and complete magnetization reversal [20].

Here we demonstrate that flatspin is able to replicate the experimental reversal processes presented in Li et al. [20], where pinwheel “diamond” ASI (Fig. 3d) is subject to an external field at different angles. A key result is that the angle  $\theta$  of the external field controls the nature of the reversal process. When  $\theta$  is small (equally aligned to both sublattices), reversal happens in a single avalanche, whereas when  $\theta$  is large (more aligned to one sublattice), reversal happens in a two-step process where one sublattice switches completely before the other. Previous attempts at capturing this behavior in a dipole model have proven unfruitful [20].

To replicate this process in flatspin, an asymmetric switching astroid is required, i.e., the threshold along the parallel component is reduced by setting  $b < 1$  in Eq. (8). From micromagnetic simulations of a single  $470 \times 170 \times 10$  nm magnet we obtain the following characteristic switching parameters:  $b = 0.28$ ,  $c = 1.0$ ,  $\beta = 4.8$  and  $\gamma = 3.0$ . Other simulation parameters include  $\alpha \approx 0.00033$ ,  $h_k = 0.098$ ,  $k_{\text{disorder}} = 0.05$  and a neighbor distance of 10.

Figs. 8a to 8d show hysteresis loops and array snapshots when the field is aligned with the array ( $\theta = 0^\circ$  and  $\theta = -6^\circ$ ). As can be seen, the results from flatspin (Figs. 8b and 8d) are qualitatively very similar to experimental results (Figs. 8a and 8c). In all cases, the ASI undergoes reversal in a single avalanche. Reversal begins at a small number of nucleation points close to the edge, followed by domain growth and domain wall movement perpendicular to the direction of the field. The simulated system appears to have an anisotropy axis of  $0^\circ$  as opposed to  $-6^\circ$  observed experimentally. Hence Fig. 8b is most similar to Fig. 8c and Fig. 8d is most similar to Fig. 8a. It should be noted that the non-zero anisotropy axis found experimentally has not yet been explained.

Figs. 8e and 8f show the hysteresis loops and array snapshots when the field is misaligned with the array ( $\theta = 30^\circ$ ). Again, flatspin simulations (Fig. 8f) replicate key features observed experimentally (Fig. 8e). Reversal now happens in two steps: the sublattice whose magnets have their easy axis most aligned with the field will switch first, followed later by the other sublattice. This two-step reversal process results in an emergent rotation of the collective magnetization. The magnetization is constrained to follow the orientation of the magnets, resulting in reversal via stripe patterns at  $45^\circ$ .

Li et al. [20] report they were unable to replicate the magnetization details using a point-dipole Monte Carlo model. One crucial difference between flatspin and their dipole model is the switching criteria. They use the simpler criteria  $\mathbf{h}_i \cdot \mathbf{m}_i < h_k^{(i)}$ , which considers only the parallel field component and will be largely inaccurate for

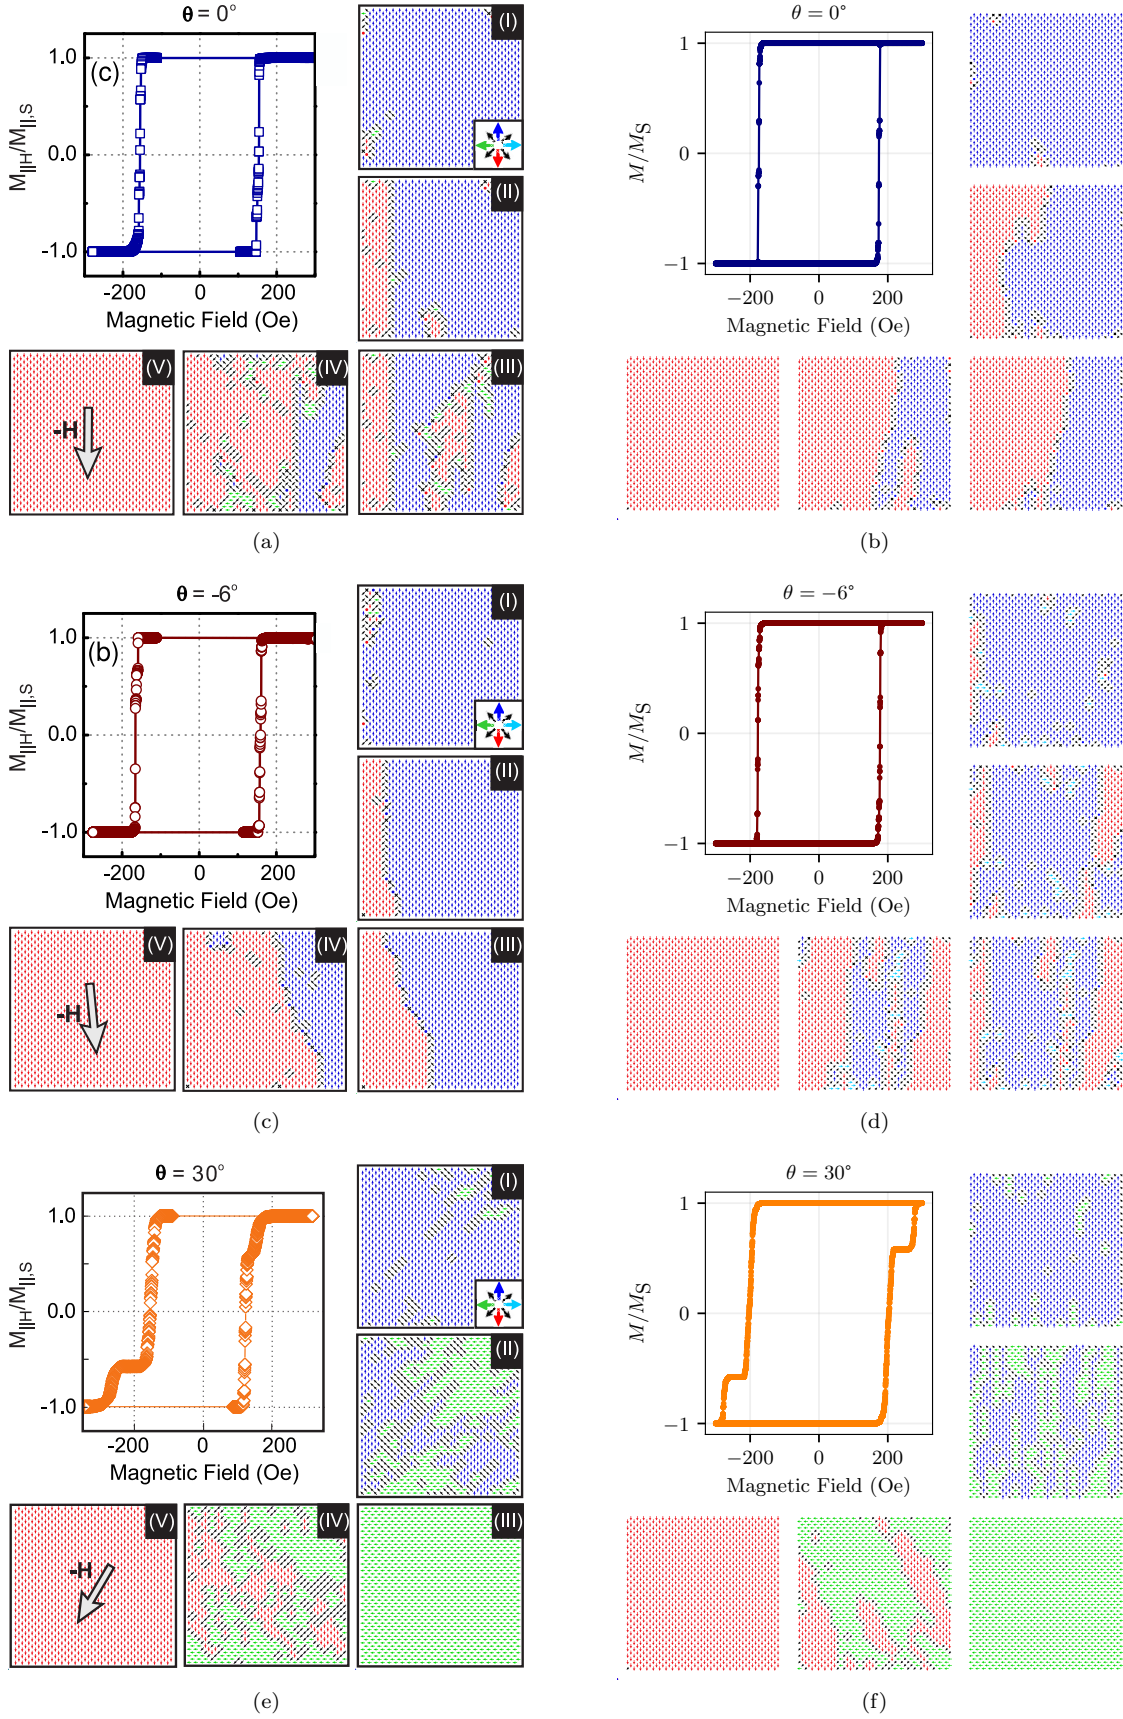


FIG. 8. Hysteresis loop and snapshots of the pinwheel units for various angles  $\theta$  of the applied field. Figures (a), (c) and (e) show experimental results, adapted from Li et al. [20], Copyright ©2018 American Chemical Society, CC-BY-4.0 <https://creativecommons.org/licenses/by/4.0/>. Figures (b), (d) and (f) show results from flatspin simulation.

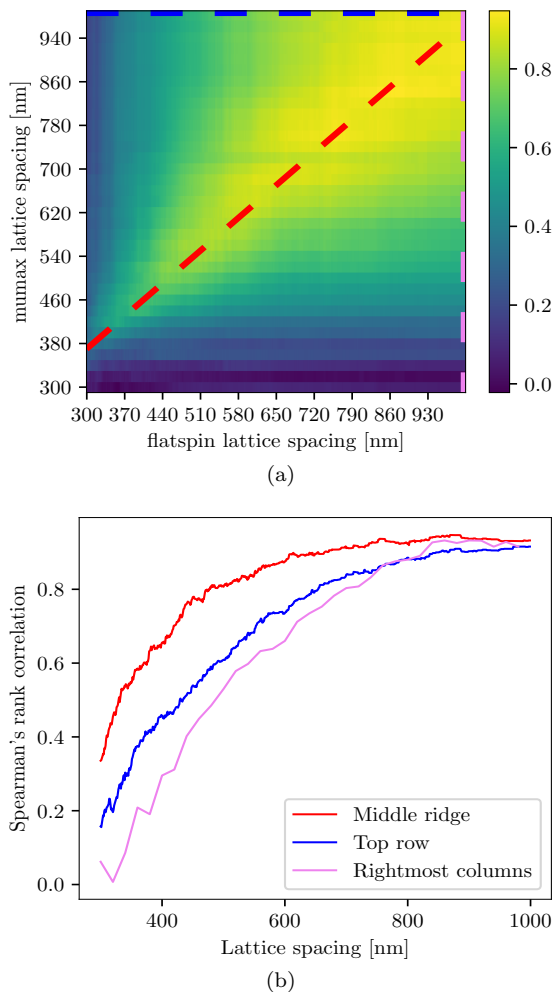


FIG. 9. (a) Spearman's rank correlation coefficients  $\rho$  averaged over 32 different square ASIs, evaluated for different lattice spacings in flatspin and MuMax3. The red line shows the approximate maximum ridge line through the heatmap. (b) The red line shows the true maximum  $\rho$  for the lattice spacing pairs. The blue and violet lines show projections of the top row and rightmost column of (a), respectively.

fields which are not aligned with the magnet's easy axis. Indeed, using this simpler switching criteria in flatspin results in a very different reversal process and magnetization details (not shown).

### E. Comparison to micromagnetic single-spin switching order

Micromagnetic simulations in MuMax3 [18, 37] have been shown to agree with experiment due to high simulation fidelity. It is, therefore, of interest to study how well flatspin agrees with MuMax3 at the level of detail expressed in flatspin.

Here we evaluate the switching strategy outlined in Section II H, by comparing the switching orders obtained in flatspin and MuMax3, of a square ASI as it undergoes

reversal by an external field. As a similarity measure, Spearman's rank correlation coefficient  $\rho$  [10] is used, where a value of 1 indicates perfect correlation and 0 indicates no correlation between switching orders.

In the weakly coupled regime, the switching order is dictated by the coercive fields of each individual magnet. In flatspin, the coercive field can be set directly by modifying  $h_k^{(i)}$ . In MuMax3, we control the coercive field implicitly, by varying the first-order, uniaxial, magnetocrystalline anisotropy,  $K_{U1}^{(i)}$  of each magnet. Given a set of randomly drawn  $K_{U1}^{(i)}$  values, the corresponding  $h_k^{(i)}$  values were obtained by a linear map.

The system we considered was a  $4 \times 4$  square (closed) ASI, each magnet measuring  $220 \text{ nm} \times 80 \text{ nm} \times 25 \text{ nm}$ . flatspin was run with parameters  $b = 0.38$ ,  $c = 1$ ,  $\beta = 1.5$ , and  $\gamma = 3.2$ . In both simulators, we applied a gradually increasing reversal field at  $\theta = 44^\circ$ .

As the lattice spacing is decreased, the dipolar interactions begin contributing to the switching order. To verify that flatspin still captures switching dynamics, we perform a comparison of the switching orders for all pairs of lattice spacings in both simulators.

Fig. 9a shows the correlations for each pair of lattice spacings as an average over 32 different square ASIs. We observe a clear linear relationship between the two simulators, with higher lattice spacings exhibiting higher correlation. The non-zero y-intercept in the heatmap indicates that, as suspected, the coupling strength is slightly underestimated by the dipole approximation employed in flatspin, in particular for lower lattice spacings. For example, flatspin with 300 nm lattice spacing is most similar to MuMax3 with 380 nm.

The red line in Fig. 9b traces the ridge in the heatmap, i.e., the highest  $\rho$ , for each flatspin lattice spacing. As can be seen, a near-perfect agreement between the simulators is found in the weakly coupled regime (high lattice spacing). As lattice spacings decrease, the effect of dipole interactions become apparent. Below 450 nm the correlation drops. Since flatspin does not account for the micromagnetic state, complete correlation is not expected.

The particular selection of  $h_k^{(i)}$  values introduces an *inherent* bias in the switching order. One might expect that this bias causes the dipole interactions to have a negligible impact on the switching order, leading to an inflated correlation between flatspin and MuMax3, regardless of lattice spacing.

The violet and blue lines, (rightmost column and top row from Fig. 9a), show the deviation from the inherent switching order in MuMax3 and flatspin, respectively. Their rapid decline confirms that the switching order is not dominated by the inherent bias for highly coupled systems. The red line clearly shows a stronger agreement between MuMax3 and flatspin (higher than the blue and violet lines), even at smaller lattice spacings.






FIG. 10. The throughput (number of field calculations per second) as a function of number of spins. Throughput is averaged over 100 simulations of each size. The test was performed on an NVIDIA Tesla V100 GPU with 32GB of RAM. Note the logarithmic scale of the axes.

## V. PERFORMANCE

Although the total simulation time will depend on many factors, it is of interest to measure how simulation time scales with the number of spins. As the number of spins are increased, simulation time will be largely dominated by the calculation of the effective field,  $\mathbf{h}_i$ , acting on each of the  $N$  spins in the lattice. Computing time for  $\mathbf{h}_{\text{dip}}^{(i)}$  depends on the number of neighbors around spin  $i$ , which is typically constant for all spins except the ones at the edges of the geometry. For large  $N$ , the number of edge magnets is negligible (in the common ASI geometries). Computing  $\mathbf{h}_i$  for all spins will take  $\mathcal{O}(N)$  time, i.e., computation time grows no faster than linear in  $N$ .

Fig. 10 shows the throughput (number of field calculations per second) as a function of the number of spins. Here a field calculation is defined as the computation of  $\mathbf{h}_i$  for a single spin  $i$ , hence for  $N$  spins there will be  $N$  such field calculations. The geometry used was square ASI (open edges) using a standard 8 spin neighborhood for calculating  $\mathbf{h}_{\text{dip}}^{(i)}$ . The throughput was averaged over 100 simulations of each size. The test was performed on an NVIDIA Tesla V100 GPU with 32GB of RAM.

At around 200 000 spins, the throughput saturates at  $10^8$  field calculations per second. On our test setup, computing  $\mathbf{h}_i$  for one million spins takes approximately 10 ms. Above 200 000 spins we are able to fully utilize the GPU resources.

To simulate the reversal of an ASI by a gradually increasing external field, at least one field calculation per spin flip is required, i.e., at least  $N$  field calculations. If the external field gradually changes with a resolution of  $K$  values, the worst case will be when all spins flip during a single field value. In this case the number of field calculations required will be  $N + K - 1$  since there will be  $K - 1$  field calculations which results in no spin flips.

The total simulation time depends largely on the par-




FIG. 11. A snapshot from flatspin simulations of  $708 \times 708$  pinwheel ASI with more than one million magnets, as it undergoes reversal by an external field. The angle of the external field is  $\theta = 0^\circ$ .

ticular experimental setup, parameters and other system characteristics. Time will be spent on things other than field calculations, e.g., organizing and writing results to storage. Hence the total simulation time will be longer than predicted by field calculations alone. As an example, the simulations from section IV D of  $25 \times 25$  pinwheel ASI with 1250 magnets took approximately 6 seconds with  $K = 2500$ , for one reversal.

Fig. 11 shows a snapshot from flatspin simulations of  $708 \times 708$  pinwheel ASI as it undergoes reversal by an external field. The ability to simulate such large systems allows an experimenter to explore phenomena at much larger scales than can be directly observed experimentally. With more than one million magnets, the simulation of array reversal took several days to complete. A video of the full reversal is available as Supplemental Material [Note1].

## VI. SQUARE ASI ROBUSTNESS TO DILUTION DEFECTS

In the following we will make use of the unique framework provided by flatspin to explore new aspects of ASI behavior.

While square ASI is perhaps one of the most investigated ASI systems, mainly due to its display of magnetically charged defects with associated Coulombic interactions and string tension [23, 33], the robustness of this system with respect to lattice defects is largely unexplored. For the Kagome ASI it has, however, been recognized that significant removal of elements in the lattice will cause the phase transition of the system to disappear [1]. Here, we use flatspin to investigate the effect of

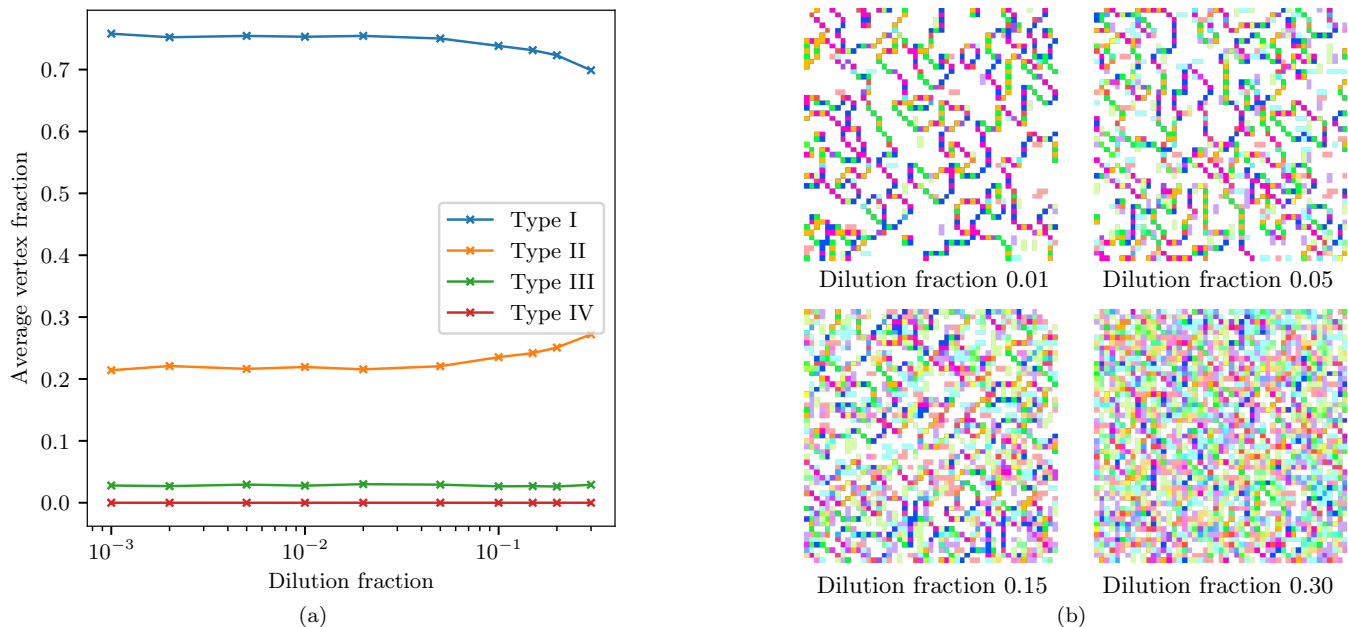


FIG. 12. (a) Average complete vertex fraction, by type, as a function of dilution by removal of spins in a  $50 \times 50$  square ASI. Only complete vertices are included. The averages are taken over 10 different instances for each dilution fraction. (b) Example end states of square ASI of different dilutions. Each pixel correspond to the net magnetization of 4 spins (vertices), as in Fig. 7b. Note that white pixels correspond to vertices of no net magnetization, i.e., type I vertices (as there are no type IV vertices).

dilution defects, i.e., random removal of lattice elements, on square vertex population and domain size in the ASI after a field demagnetization protocol.

Large ASI systems are needed to observe large-scale phenomena such as domain formation and the effects of sparse random dilution. Random defects result in stochastic behavior, as the results will depend on the exact configuration of the defects. Hence, to provide sufficient statistics of these phenomena, the experiments must be repeated for different random configurations. Furthermore, this is combined with long field protocols which are required to relax each system to a lower energy configuration.

The two-fold degenerate ground state configurations in square ASI consist of antiferromagnetic ordering on the vertical and horizontal sublattices [39]. Reaching this ground-state configuration in experimental realization has proven difficult. Relaxation through thermal annealing or field protocols will typically result in domains of charge neutral type I vertices separated by domain walls comprised of mostly type II vertices with a net magnetization (and no net magnetic charge). These domain walls, strings of type II vertices, are necessarily terminated by type III (or less frequently: type IV) vertices of opposite net magnetic charge, or at incomplete vertices, such as at the edges of the ensemble.

Diluting the ASI lattice by removal of elements introduces more vertices with uncompensated magnetic charge. Here, we use flatspin to explore how such doping with fixed magnetic charges affects the magnetic vertex

population of a square ASI.

A  $50 \times 50$  square ASI is initialized uniformly and subjected to a rotating magnetic field with decreasing amplitude. The initial amplitude is slightly larger than a saturating field, and is decreased linearly over 1000 periods. The final periods do not alter the state of the ASI. Individual spins are assigned a  $h_k = 0.2$  with 5% disorder, and the field amplitude, directed at  $45^\circ$ , starts at  $h = 0.080$  and ends at  $h = 0.072$ . Each spin is again modeled as a rectangular magnet, with astroid parameters  $b = 0.38$ ,  $c = 1.0$ ,  $\beta = 1.3$ ,  $\gamma = 3.6$ , and  $\alpha = 30272$ , with a lattice spacing  $a = 300$  nm. A neighbor distance of 3 lattice constants is used in the simulations. Larger neighbor distances (up to 50 lattice constants) were sampled for the undiluted system and revealed no discrepancy from a neighbor distance of 3. Spins in the lattice are removed randomly, until the desired dilution fraction is achieved. Ten different instances of this random removal were performed for each level of dilution.

After the field protocol described above, the *complete* (undiluted) vertices of the final state were counted and classified according to their vertex type. The resulting vertex fraction count of complete vertices can be seen in Fig. 12a.

Somewhat surprisingly, the fraction of each vertex type at the end of the field protocol is almost constant through dilution. We also note that a large fraction of the complete vertices ( $> 0.70$ ) obey the ice rules (type I or type II), indicating that the system is strongly coupled.

Another apparent effect of dilution is a dramatic reduc-



tion in domain size, see Fig. 12b. It is interesting to note that, despite the change of domain size, which should include more domain boundaries and more magnetized vertices, the type population fraction stays constant through dilution. Our findings indicate that the fixed magnetic charges introduced through dilution comes in addition to the mobile magnetic charges. While the domain walls in systems with low dilution are comprised of ordinary type II vertices, the fixed incomplete vertices add up with complete type II vertices to form more domain walls for the systems with higher dilution and, thus, smaller domains.

The presented simulation study demonstrates that the type population in square ASI is robust towards dilution. However, the sizes of ground state domains reduce considerably as an increasing number of magnetic islands are removed at random, as seen in [39]. This preliminary study into the effect of dilution defects in square ASI demonstrates how flatspin can be used as an efficient tool to explore new physical phenomena in ASI systems, which in turn can be verified by experiments.

## VII. CONCLUSION

flatspin is a highly effective simulator for ASI ensemble dynamics. At its heart lies a robust magnetic model based on dipole-dipole interactions with a switching criteria based on a generalized Stoner-Wohlfarth model. Accompanying the model is a toolbox of useful input encoders and analysis tools. The model includes several common ASI geometries, and there are no inherent limits to the range of possible geometries.

The flatspin ASI model has been verified against mi-

cro-magnetic simulations and experimental results from the literature. Emergent fine-scale patterns in kagome ASI were replicated successfully, where the formation of Dirac strings matched experimental results. Large-scale domain sizes in square ASI were reproduced, and good agreement was found between flatspin and experimental results. Using flatspin, the experimental magnetization reversal of pinwheel ASI was reproduced for the first time in a dipole model. On a detailed level, we found good agreement between micromagnetic simulations and flatspin in terms of magnet switching order.

Finally, we shed light on the effects of dilution defects in square ASI. Our investigation revealed a surprisingly robust vertex type population, as random magnets were removed from the lattice.

Through GPU acceleration, flatspin scales to large ASI systems with millions of magnets. High speed, parallel computation allows for a large number of ASI simulations to be executed, enabling quick exploration of parameters and novel geometries. The flexibility and performance offered by flatspin opens for unprecedented possibilities in ASI research.

## ACKNOWLEDGMENTS

We thank Jonathan Leliaert for fruitful discussions on the theory of magnetic modeling. This work was funded in part by the Norwegian Research Council by the IK-TPLUSS project SOCRATES (Grant no. 270961), and in part by the EU FET-Open RIA project SpinENGINE (Grant no. 861618). Simulations were executed on the NTNU EPIC compute cluster [32].

- 
- [1] Andriushchenko, P. (2019). Influence of cutoff dipole interaction radius and dilution on phase transition in kagome artificial spin ice.
  - [2] Anghinolfi, L., Luetkens, H., Perron, J., Flokstra, M. G., Sendetskyi, O., Suter, A., Prokscha, T., Derlet, P. M., Lee, S. L., and Heyderman, L. J. (2015). Thermodynamic phase transitions in a frustrated magnetic metamaterial. *Nature Communications*, 6(1):8278.
  - [3] Bar-Yam, Y. (1997). *Dynamics of Complex Systems*. Perseus Books, USA.
  - [4] Brown, W. F. (1963). Thermal fluctuations of a single-domain particle. *Phys. Rev.*, 130:1677–1686.
  - [5] Budrikis, Z. (2012). *Athermal dynamics of artificial spin ice: disorder, edge and field protocol effects*. PhD thesis, The University of Western Australia.
  - [6] Fan, C. and Wu, F. Y. (1969). Ising model with second-neighbor interaction. i. some exact results and an approximate solution. *Phys. Rev.*, 179:560–569.
  - [7] flatspin contributors (2020a). flatspin user manual. <https://flatspin.readthedocs.io/en/latest/>.
  - [8] flatspin contributors (2020b). flatspin website. <https://flatspin.gitlab.io>.
  - [9] Frenkel, J. and Doefman, J. (1930). Spontaneous and induced magnetisation in ferromagnetic bodies. *Nature*, 126(3173):274–275.
  - [10] Gibbons, J. D. and Chakraborti, S. (2014). *Nonparametric Statistical Inference: Revised and Expanded*. CRC press.
  - [11] Gillette, P. R. and Oshima, K. (1958). Magnetization reversal by rotation. *Journal of Applied Physics*, 29(3):529–531.
  - [12] Heyderman, L. J. and Stamps, R. L. (2013). Artificial ferroic systems: novel functionality from structure, interactions and dynamics. *Journal of Physics: Condensed Matter*, 25(36):363201.
  - [13] Jensen, J. H., Folven, E., and Tufte, G. (2018). Computation in artificial spin ice. In *The 2018 Conference on Artificial Life*, pages 15–22, Tokyo, Japan. MIT Press.
  - [14] Jensen, J. H. and Tufte, G. (2020). Reservoir computing in artificial spin ice. *Artificial Life Conference Proceedings*, (32):376–383.
  - [15] Ke, X., Li, J., Nisoli, C., Lammert, P. E., McConville, W., Wang, R. F., Crespi, V. H., and Schiffer, P. (2008). Energy Minimization and ac Demagnetization in a Nanomagnet Array. *Physical Review Letters*, 101(3):037205.
  - [16] Kikuchi, R. (1956). On the minimum of magnetization

- reversal time. *Journal of Applied Physics*, 27(11):1352–1357.
- [17] Kittel, C. (1946). Theory of the structure of ferromagnetic domains in films and small particles. *Phys. Rev.*, 70:965–971.
- [18] Leliaert, J., Dvornik, M., Mulkers, J., De Clercq, J., Milošević, M. V., and Van Waeyenberge, B. (2018). Fast micromagnetic simulations on GPU - Recent advances made with mumax3. *Journal of Physics D: Applied Physics*, 51(12).
- [19] Levis, D., Cugliandolo, L. F., Foini, L., and Tarzia, M. (2013). Thermal Phase Transitions in Artificial Spin Ice. *Physical Review Letters*, 110(20):207206.
- [20] Li, Y., Paterson, G. W., Macauley, G. M., Nascimento, F. S., Ferguson, C., Morley, S. A., Rosamond, M. C., Linfield, E. H., MacLaren, D. A., Macêdo, R., Marrows, C. H., McVitie, S., and Stamps, R. L. (2019). Superferromagnetism and Domain-Wall Topologies in Artificial “Pinwheel” Spin Ice. *ACS Nano*, page acsnano.8b08884.
- [21] Macêdo, R., Macauley, G. M., Nascimento, F. S., and Stamps, R. L. (2018). Apparent ferromagnetism in the pinwheel artificial spin ice. *Physical Review B*, 98(1):014437.
- [22] Mengotti, E., Heyderman, L. J., Rodríguez, A. F., Noltling, F., Hügli, R. V., and Braun, H.-B. (2011). Real-space observation of emergent magnetic monopoles and associated dirac strings in artificial kagome spin ice. *Nature Physics*, 7(1):68.
- [23] Mól, L. A., Silva, R. L., Silva, R. C., Pereira, A. R., Moura-Melo, W. A., Costa, B. V., and Carlo, M. (2009). Magnetic monopole and string excitations in two-dimensional spin ice ARTICLES YOU MAY BE INTERESTED IN. *J. Appl. Phys.*, 106:63913.
- [24] Morris, D. J. P., Tennant, D. A., Grigera, S. A., Klemke, B., Castelnovo, C., Moessner, R., Czternasty, C., Meissner, M., Rule, K. C., Hoffmann, J.-U., Kiefer, K., Gerischer, S., Slobinsky, D., and Perry, R. S. (2009). Dirac strings and magnetic monopoles in the spin ice Dy<sub>2</sub>Ti<sub>2</sub>O<sub>7</sub>. *Science*, 326(5951):411–414.
- [25] Néel, L. (1949). Théorie du traînage magnétique des ferromagnétiques en grains fins avec applications aux terres cuites. *Ann. géophys.*, 5:99–136.
- [26] Newman, M. E. J. and Barkema, G. T. (1999). *Monte Carlo Methods in Statistical Physics*. Clarendon Press.
- [Note1] Note1. See Supplemental Material at [URL will be inserted by publisher] for a video of the reversal of 708x708 pinwheel ASI.
- [28] Qi, Y., Brintlinger, T., and Cumings, J. (2008). Direct observation of the ice rule in an artificial kagome spin ice. *Physical Review B*, 77(9):094418.
- [29] Rougemaille, N., Moutaigne, F., Canals, B., Duluard, A., Lacour, D., Hehn, M., Belkhou, R., Fruchart, O., El Moussaoui, S., Bendouan, A., et al. (2011). Artificial kagome arrays of nanomagnets: a frozen dipolar spin ice. *Physical Review Letters*, 106(5):057209.
- [30] Saccone, M., Scholl, A., Velten, S., Dhuey, S., Hofhuis, K., Wuth, C., Huang, Y.-L., Chen, Z., Chopdekar, R. V., and Farhan, A. (2019). Towards artificial ising spin glasses: Thermal ordering in randomized arrays of ising-type nanomagnets. *Phys. Rev. B*, 99:224403.
- [31] Sinwani, O., Reiner, J. W., and Klein, L. (2014). Monitoring superparamagnetic Langevin behavior of individual SrRuO<sub>3</sub> nanostructures. *Physical Review B - Condensed Matter and Materials Physics*, 89(2):1–5.
- [32] Sjölander, M., Jahre, M., Tufte, G., and Reissmann, N. (2019). EPIC: An energy-efficient, high-performance GPGPU computing research infrastructure. arXiv:1912.05848 [cs.DC].
- [33] Skjærvø, S. H., Marrows, C. H., Stamps, R. L., and Heyderman, L. J. (2020). Advances in artificial spin ice. *Nature Reviews Physics*, 2(1):13–28.
- [34] Sklenar, J., Lao, Y., Albrecht, A., Watts, J. D., Nisoli, C., Chern, G.-W., and Schiffer, P. (2019). Field-induced phase coexistence in an artificial spin ice. *Nature Physics*, 15(2):191–195.
- [35] Tanaka, M., Saitoh, E., Miyajima, H., Yamaoka, T., and Iye, Y. (2006). Magnetic interactions in a ferromagnetic honeycomb nanoscale network. *Physical Review B*, 73(5):052411.
- [36] Tannous, C. and Gieraltowski, J. (2008). The stoner–wohlfarth model of ferromagnetism. *European Journal of Physics*, 29(3):475–487.
- [37] Vansteenkiste, A., Leliaert, J., Dvornik, M., Helsen, M., Garcia-Sanchez, F., and Van Waeyenberge, B. (2014). The design and verification of MuMax3. *AIP Advances*, 4(10):107133.
- [38] Wang, R. F., Nisoli, C., Freitas, R. S., Li, J., McConville, W., Cooley, B. J., Lund, M. S., Samarth, N., Leighton, C., Crespi, V. H., and Schiffer, P. (2006). Artificial ‘spin ice’ in a geometrically frustrated lattice of nanoscale ferromagnetic islands. *Nature*, 439(7074):303–306.
- [39] Zhang, S., Gilbert, I., Nisoli, C., Chern, G. W., Erickson, M. J., O’Brien, L., Leighton, C., Lammert, P. E., Crespi, V. H., and Schiffer, P. (2013). Crystallites of magnetic charges in artificial spin ice. *Nature*, 500(7464):553–557.