

flatspin: A large-scale artificial spin ice simulator

Johannes H. Jensen ^{1,*}, Anders Strømberg ^{1,*}, Odd Rune Lykkebø ¹, Arthur Penty ¹, Jonathan Leliaert ²,
Magnus Sjölander ¹, Erik Folven ¹ and Gunnar Tufte ¹

¹Norwegian University of Science and Technology, 7034 Trondheim, Norway

²Department of Solid State Sciences, Ghent University, 9000 Ghent, Belgium



(Received 17 February 2020; revised 24 June 2022; accepted 24 June 2022; published 4 August 2022)

We present flatspin, a novel simulator for systems of interacting mesoscopic spins on a lattice, also known as artificial spin ice (ASI). A generalization of the Stoner-Wohlfarth model is introduced, and combined with a well-defined switching protocol to capture realistic ASI dynamics using a point-dipole approximation. Temperature is modelled as an effective thermal field, based on the Arrhenius-Néel equation. Through GPU acceleration, flatspin can simulate the dynamics of millions of magnets within practical time frames, enabling exploration of large-scale emergent phenomena at unprecedented speeds. We demonstrate flatspin's versatility through the reproduction of a diverse set of established experimental results from literature. In particular, the field-driven magnetization reversal of "pinwheel" ASI is reproduced, for the first time, in a dipole model. Finally, we use flatspin to explore aspects of "square" ASI by introducing dilution defects and measuring the effect on the vertex population.

DOI: [10.1103/PhysRevB.106.064408](https://doi.org/10.1103/PhysRevB.106.064408)

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 [1,2]. Using modern nanofabrication techniques, emergent phenomena can be facilitated through direct control of the ASI geometry, e.g., collective ferromagnetic/antiferromagnetic ordering [3], Dirac strings [4], and phase transitions [5,6]. 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 [7,8]. Computer simulations have proven invaluable to gain insight into the rich behavior of these coupled systems.

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 macrospin

approximated by a point dipole [9]. Conventionally, Monte Carlo methods have been used in conjunction with the dipole approximation to search for low energy configurations [10,11] or to study statistical measures such as vertex populations [9]. However, Monte Carlo methods do not model the dynamical pathway taking the system from an initial configuration to the final low-energy configuration. They are inherently stochastic and better suited for ensemble statistics rather than dynamics [12].

Here we present flatspin, a point-dipole simulator for large ASI systems that is capable of capturing realistic dynamics, at long time scales. We introduce a generalization of the Stoner-Wohlfarth model, that describes the switching characteristics of numerous nanomagnet shapes. This generalized model is combined with a novel, well-defined, switching protocol, to capture the dynamics of large ASI systems. All influences on the magnets are represented by magnetic fields, including a stochastic thermal field derived from the Arrhenius-Néel equation. These crucial aspects of flatspin, combined with GPU acceleration, extends the possibilities of ASI simulation. Using flatspin in place of micromagnetic simulations increases the possible simulation sizes from hundreds to millions of magnets, enabling exploration of large-scale, emergent phenomena in these intriguing systems.

In this paper, we present the motivation and design of flatspin. The thermal model is verified against established analytical and numerical models. 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 [13]. 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.

*These authors contributed equally to this work.

†Corresponding author: johannes.jensen@ntnu.no

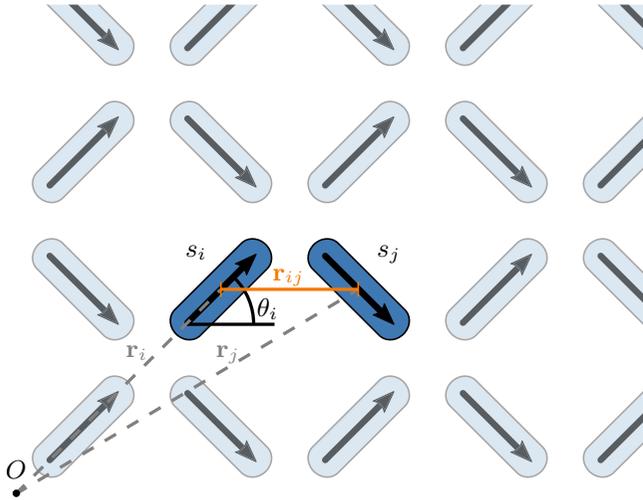


FIG. 1. The representation of nanomagnets as spins s_i and associated quantities: position \mathbf{r}_i , angle θ_i , and distance to neighbor j , \mathbf{r}_{ij} . Note that the magnetization of spin i is given by its spin s_i , and orientation θ_i .

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 (Sec. II A), and each dipole can be affected by three types of external influence: magnetic dipole-dipole coupling (Sec. II C), an applied external magnetic field (Sec. II D), and thermal fluctuations (Sec. II E). The switching of spins is determined using a generalized Stoner-Wohlfarth model, which takes the shape anisotropy of the simulated nanoislands into account (Sec. II F). Imperfections in the ASI are introduced as different coercive fields, set per spin (Sec. II G). Dynamics are modeled using a deterministic single spin flip strategy (Sec. 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 to exhibit a single ferromagnetic domain, i.e., coherent magnetization throughout the magnet. The single domain state is stable as the energy cost associated with domain walls exceeds the cost associated with the demagnetization energy [14,15]. 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 shape anisotropy of the thin elongated islands will restrict their magnetization to two possible in-plane directions. Hence, individual magnets can be approximated by classical macrospins 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 binary macrospin, $s_i \in \{-1, +1\}$.

As illustrated in Fig. 1, each magnetic dipole is modelled with a position \mathbf{r}_i and rotation θ_i , which together define the ASI geometry. Furthermore, each magnet is assigned a coercive field, $h_c^{(i)}$, describing its resistance to switching (see Sec. II F). 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 $\hat{\mathbf{m}}_i = [\cos \theta_i, \sin \theta_i]$ 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 stochastic magnetic field representing thermal fluctuations in each magnetic element. Each of these field contributions are 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 α scales the dipolar coupling strength between spins. The coupling strength α is given by $\alpha = \frac{\mu_0 M}{4\pi a^3}$, where a is the lattice spacing, M is the net magnetic moment of a single magnet, and μ_0 is the vacuum permeability. The distance \mathbf{r}_{ij} is thus 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}_{dip} [16,17], 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 [18]. The flatspin documentation [19] provides an example of how to choose an appropriate neighborhood distance.

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}_{ext}), or as a spatial vector field ($\mathbf{h}_{\text{ext}}(\mathbf{r})$).

Time-dependent external fields are supported, i.e., \mathbf{h}_{ext} is a discrete time series of either local, global, or spatial fields. A variety of time-dependent external fields come predefined with flatspin, 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 flatspin, the thermal energy fluctuations of individual magnets $E_{\text{th}}^{(i)}$ are represented by a corresponding stochastic magnetic field $\mathbf{h}_{\text{th}}^{(i)}$. The following section describes how the thermal field magnitude is derived.

In a physical ensemble of particles, there is a thermal energy budget, on the scale of $k_{\text{B}}T$, where k_{B} and T denote the Boltzmann constant and the temperature, respectively. For bistable magnetic particles, the thermal energy causes random switching events, at a characteristic rate given by the Arrhenius-Néel equation [20],

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

where f_0 is the attempt frequency and ΔE is the particle's energy barrier for switching.

The energy barrier ΔE corresponds to the *additional* Zeeman energy required for magnetization reversal. This additional energy is a function of the smallest additional field needed for switching Δh . The Zeeman energy from Δh is given by $\Delta E = \Delta h M_{\text{th}}$, where M_{th} is the thermal nucleation moment. Note that M_{th} is typically smaller than the entire magnetic moment M , since thermal nanomagnetic switching is generally mediated by a smaller nucleation volume. The relevant criterion for switching is the magnitude of the stochastic thermal field compared to the minimum energy barrier.

The probability of thermal switching follows a Poisson distribution $P_r(k, \Delta t, f)$, where k is the number of switching events in a time interval Δt , and f is the characteristic switching rate given by Eq. (4). The probability of switching at least once P_{switch} is given by

$$\begin{aligned} P_{\text{switch}} &= P_r(k > 0), \\ &= 1 - P_r(k = 0), \\ &= 1 - \exp(-f \Delta t), \\ &= 1 - \exp\left(\Delta t f_0 \exp\left(-\frac{\Delta h M_{\text{th}}}{k_{\text{B}}T}\right)\right). \end{aligned} \quad (5)$$

In our model, we use this expression to approximate the probability that a magnet will switch *once*. For this to be valid, Δt must be sufficiently small so that multiple switching events are unlikely. In other words, it is assumed that $P_r(k > 1) \ll P_r(k = 1)$, so that magnets do not have time to “switch back” during the time Δt . This is relevant only for weak bias fields, where the probability of switching is nearly symmetrical. With any significant bias field, the probability of “switching back” will be negligible.

It is important to note that flatspin does not account for the temperature dependence of the material parameters. If

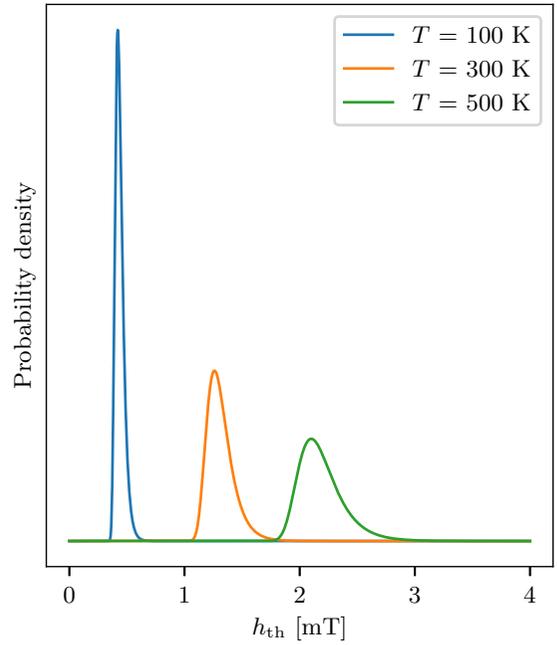


FIG. 2. Probability density functions (PDF) for h_{th} at different temperatures T . The example plots are for a magnetic particle with saturation magnetization $M_{\text{S}} = 860 \text{ kA m}^{-1}$ and volume $V = 220 \text{ nm} \times 80 \text{ nm} \times 3 \text{ nm}$, for a time interval $\Delta t = 1 \text{ ms}$.

these parameters are expected to vary significantly in the temperature range of interest, e.g., M_{th} , this has to be explicitly accounted for by the user.

As discussed in Sec. II B, all magnetic influences are included as magnetic fields, and their sum is compared to the switching condition to determine whether a magnet switches. Within this framework, we now derive an expression for the thermal field based on P_{switch} .

Consider a stochastic field variable X from which a thermal field h_{th} is sampled. The probability of drawing a thermal field h_{th} larger than the minimum switching field Δh equals the switching probability for the same minimum switching field,

$$P(X > \Delta h) = P_{\text{switch}}(\Delta h), \quad (6)$$

$$1 - P(X \leq \Delta h) = 1 - \exp(-f \Delta t), \quad (7)$$

$$P(X \leq \Delta h) = \exp(-f \Delta t). \quad (8)$$

$P(X \leq \Delta h)$ is the cumulative density function (CDF) of the distribution for h_{th} that matches the Poisson distribution in Eq. (5). Using inverse transform sampling, we use the expression for the CDF to transform a uniformly distributed random number u to a thermal magnetic field magnitude h_{th} ,

$$h_{\text{th}} = \frac{-k_{\text{B}}T}{M_{\text{th}}} \ln\left(\frac{\ln(u)}{-\Delta t f_0}\right). \quad (9)$$

Figure 2 illustrates how temperature influences the probability density function for the stochastic thermal field. As can be seen, both the expected value and the variance of h_{th} increases with temperature. In other words, the magnitude and spread of the stochastic field h_{th} increases, effectively increasing the probability of thermal flips.

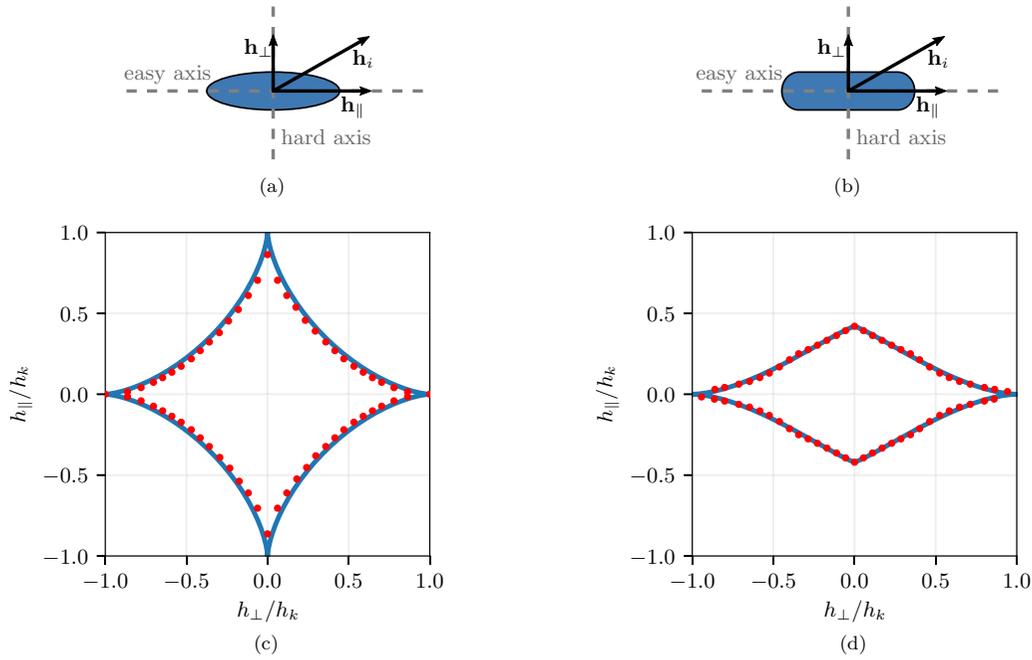


FIG. 3. Top: Schematic showing hard and easy axes of (a) an elliptical magnet and (b) a rectangular stadium-shaped magnet, as well as the total field acting on the magnet \mathbf{h}_i with its parallel and perpendicular components, \mathbf{h}_\parallel and \mathbf{h}_\perp , respectively. Bottom: Switching 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. (11). The astroids have been normalized with respect to h_k .

A stochastic thermal field magnitude drawn from h_{th} is converted to a vector field \mathbf{h}_{th} parallel to $\Delta\mathbf{h}$ (the smallest additional field for switching). When the thermal field is added to the sum of fields for each magnet, the probability of switching will follow the Poisson distribution in Eq. (5). In this way, thermal fluctuations are modeled as an additional local field $\mathbf{h}_{\text{th}}^{(i)}$ applied to each magnet individually.

F. Switching

Magnetization reversal, or *switching*, may take place when a magnet is subjected to a magnetic field or as a result of thermal fluctuations. If the field is sufficiently strong and directed against the magnetization \mathbf{m}_i , the magnetization will switch direction.

The critical field strength for switching 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. 3(a), 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. For any applied field \mathbf{h}_i that is outside the switching astroid, the magnet will switch as long as the field is directed against the current magnetization.

Figure 3(c) shows the normalized switching astroid for an elliptical magnet [Fig. 3(a)] as obtained from micromagnetic

simulations using MuMax3 [21]. 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 [22]. The characteristic SW astroid is shown in Fig. 3(c) (blue line) and is described by the equation

$$\left(\frac{h_\parallel}{h_k}\right)^{2/3} + \left(\frac{h_\perp}{h_k}\right)^{2/3} = 1. \quad (10)$$

In the SW model, switching may occur when the left-hand side of Eq. (10) is greater than one.

The astroid obtained from micromagnetic simulations and the SW astroid [Fig. 3(c)] are nearly identical. Despite its simplicity, the SW model clearly captures the switching behavior of elliptical nanomagnets.

However, the SW model is only accurate for elliptical magnets. Other magnet shapes typically have quite different switching characteristics. Figure 3(d) shows the switching astroid for rectangular stadium-shaped magnets (red dots), which is the shape commonly used in most fabricated ASIs [Fig. 3(b)]. Notice how the astroid is *asymmetric*: Rectangular magnets switch more easily with a field applied along the easy axis than the hard axis.

To capture the asymmetric switching characteristics of nonelliptical magnets, we have generalized the SW switching model to allow an asymmetry between easy and hard axes. Additionally, the model has been extended to allow for tuning of the curvature of the extrema. In the generalized model, the

switching threshold is given by

$$\left(\frac{h_{\parallel}}{bh_k}\right)^{2/\gamma} + \left(\frac{h_{\perp}}{ch_k}\right)^{2/\beta} = 1, \quad (11)$$

where b , c , β , and γ are parameters, which adjust the shape of the astroid: b and c define the height and width, respectively, while β and γ 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. (11) reduces to Eq. (10), 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. 3(d) (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. (11) 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 ,

$$\mathbf{h}_i \cdot \mathbf{m}_i < 0. \quad (12)$$

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_{disorder} defines the distribution of coercive fields, i.e., $h_k^{(i)}$ is sampled from a normal distribution $\mathcal{N}(h_k, \sigma)$, where $\sigma = k_{\text{disorder}} \cdot h_k$. Negative $h_k^{(i)}$ values are disallowed.

H. Dynamics

Given the total magnetic field acting on each spin, \mathbf{h}_i , flatspin employs deterministic single spin flip dynamics. At each simulation step, we calculate \mathbf{h}_i , which will contain a stochastic term in the case of nonzero temperature. Next, we determine which spins *may* flip according to the switching criteria Eqs. (11) and (12). 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. (11) is the greatest. Ties are broken in a deterministic, arbitrary manner, although with nonzero 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 quasistatic compared to the timescale 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 subnanosecond range [23,24], and experimental setups often employ external magnetic fields, which can be considered static at this timescale. At high applied field frequencies, more complex physical phenomena such as spin waves will have a non-negligible effect on switching dynamics. Such high-frequency phenomena are not considered in flatspin.

Assumption 2 holds if the coercive fields $h_c^{(i)}$, and total field \mathbf{h}_i , of the magnets are sufficiently nonuniform, so that there will always be a single magnet that will flip first. It is assumed to be unlikely that two magnets will have a $h_c^{(i)}$ and \mathbf{h}_i that bring them equally far outside the switching astroid. However, in those rare cases where two magnets are equally far outside, 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 *quasistatic 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 quasistatic 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 geometries have been proposed in the literature. Figure 4 depicts the geometries included in flatspin, which are the most commonly used ASI geometries: square [16], kagome [25,26], pinwheel [13,18], and Ising [27]. Note that when we refer to "pinwheel ASI" in this paper, we are referring explicitly to the 45° variant.

Geometries are often decomposed into two or more "sublattices", where the magnets within one sublattice are all aligned, i.e., have the same rotation. In Fig. 4, 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 θ_i .

J. Limitations of the model

The flatspin model makes several assumptions and approximations, which means there are inherent limitations to what

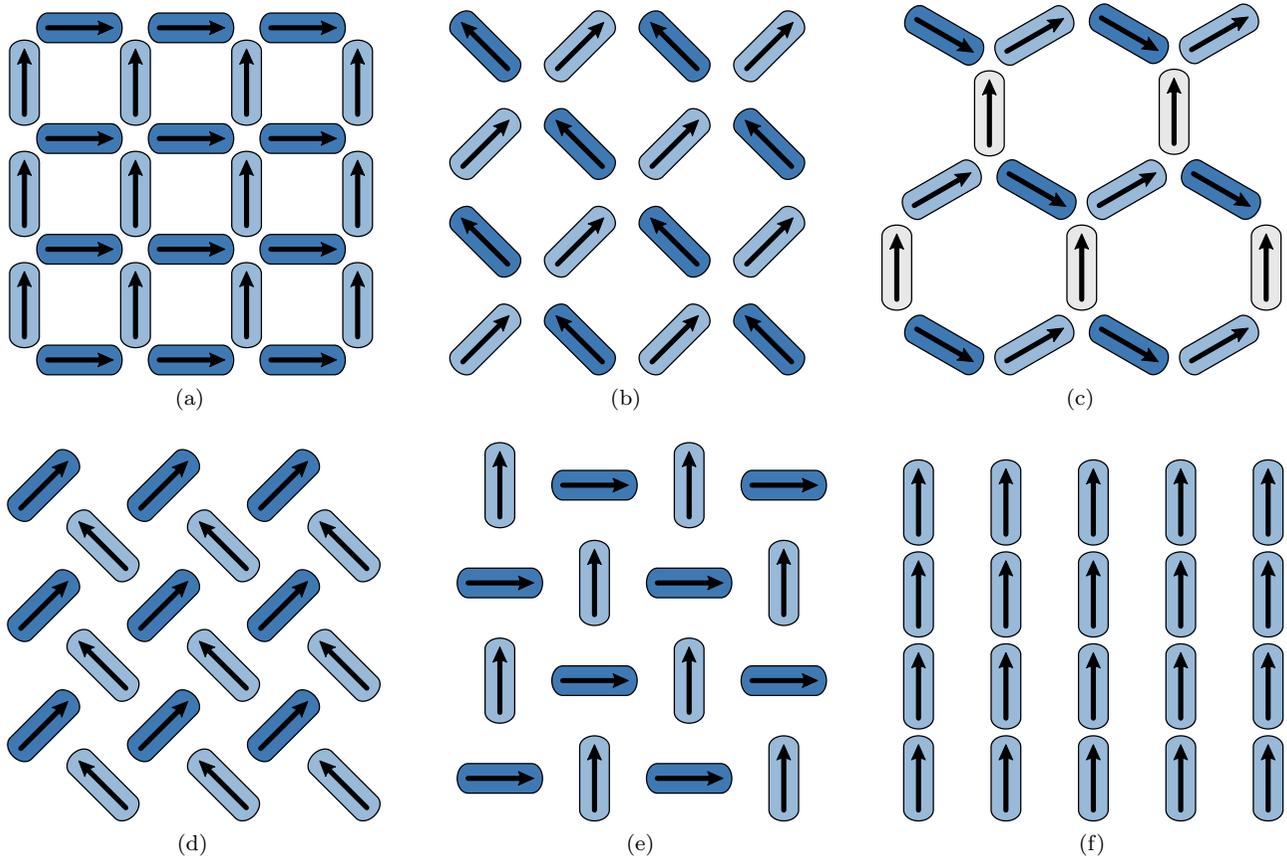


FIG. 4. flatspin includes the most common ASI geometries: (a) Square (closed edges), (b) square (open edges), (c) kagome, (d) pinwheel “diamond”, (e) pinwheel “lucky-knot”, and (f) Ising.

physics flatspin can capture. In this section, we outline the main limitations of the model.

All magnets in the model are approximated as binary mesoscopic spins, i.e., the magnetization direction of a single magnet is always parallel to the easy axis of the magnet. In reality, the magnetization of nanomagnet islands may deviate somewhat from the easy axis, which in turn would influence the dipolar fields.

It is assumed that the size and shape of all magnets is identical, as all magnets have the same net magnetic moment M , and the same switching astroid. Magnet imperfections are modelled solely as a disorder in the coercive fields, i.e., without any effect on the magnetic moment. Hybrid systems with magnets of different size and shape are therefore not supported.

The point dipole approximation underestimates the coupling coefficients for small lattice spacings. As the lateral dimensions of the magnets are not taken into account, the physical proximity of the magnets is underestimated. This can be remedied by artificially increasing the coupling strength α for highly coupled systems.

The dynamical model assumes switching to be instantaneous. In reality, magnetic switching takes finite time, mediated by a rotation of the internal magnetization state. Such transient states may affect ensemble dynamics in subtle ways, which will not be captured in the instantaneous model.

Another limitation of the dynamical model is that switching is assumed to be sequential. While simultaneous switching is possible in reality, it is not modeled in flatspin.

In spite of these limitations, the flatspin model is able to capture a range of real-world phenomena, as we will see in Sec. V.

While flatspin is specifically designed for artificial spin ice consisting of ferromagnetic macrospins, the model could be modified to accommodate other forms of artificial spin systems with bi-stable elements. Some examples include colloidal spin ice [28,29], macroscopic magnets [30], interacting skyrmions [31], and superconducting vortices [32]. Simulating such systems would require three main changes to the flatspin magnetic model. First, the magnetic dipole-dipole interactions (Sec. II C) would need to be replaced with a suitable interaction field along with the interaction modifier α . Second, the thermal field (Sec. II E) should be modified to include any other effects of temperature in the relevant system. Finally, a suitable switching mechanism (Sec. II F) must be devised to capture the switching barriers as a function of the total fields.

K. A note on units

The physical unit of the \mathbf{h} field in flatspin is Tesla [T]. While the \mathbf{H} field is typically described in units of [A m^{-1}], the fields in flatspin are exclusively external to the magnets. In

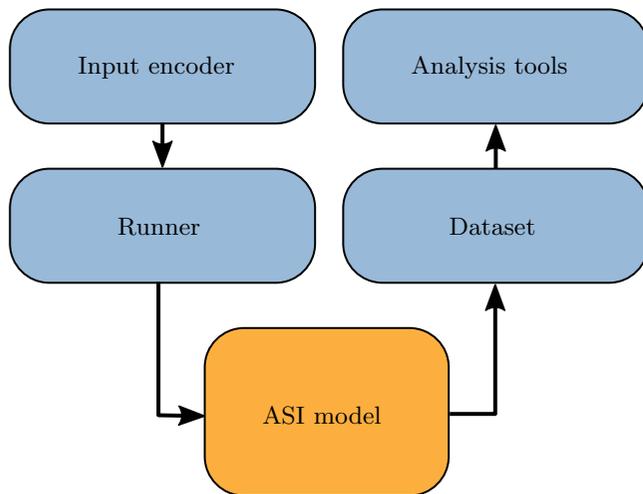


FIG. 5. Overview of the flatspin architecture, with arrows indicating data flow.

other words, the \mathbf{h} field is equivalent to a \mathbf{B} field in the absence of material magnetization, i.e., $\mathbf{h} = \mu_0 \mathbf{H}$. Correspondingly, the magnetic moments M and M_{th} have units $[\text{A m}^2]$.

III. SIMULATION FRAMEWORK

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

Figure 5 illustrates the overall architecture of flatspin. The *ASI model* has been described in detail in Sec. II. Conceptually, the ASI model describes the physical system under study. The rest of the components are tools used by the researcher 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. 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 computing cluster.

Results are stored in a well-defined *dataset* format, which makes the analysis of a large number 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 [19].

IV. VERIFICATION OF FLATSPIN

The flatspin software has been verified through an extensive suite of unit tests, where computed results are compared to theoretical values. We do not go into detail about the unit tests here, but the test suite is packaged with the flatspin software, and available from the website [19].

While the unit tests verify the software implementation, a verification of the temperature model itself (Sec. II E) is necessary. In the next section, we compare flatspin simulations to experiments where the results are known analytically.

Stochastic thermal field

To verify the temperature model in flatspin, we consider the effect of temperature on a system of noninteracting spins at equilibrium. In particular, we investigate the relationship between the temperature scale and the thermal fields, described in Sec. II E. In the following, the magnetization behavior of magnets subjected to an external field and temperature, $M(H, T)$, is simulated and compared to results of other established techniques, both analytical and numerical.

Two different scenarios are considered:

- (1) The coercive fields are small compared to the external field, and switching is mostly an effect of the external field competing with temperature.
- (2) The switching threshold is comparable to the external field and switching is also influenced by the energy landscape of the magnet, as captured by the shape of the switching astroid.

For scenario (1) we use a switching threshold of $h_k = 1$ mT, and for scenario (2) we use $h_k = 20$ mT. For both scenarios, we use the unaltered Stoner-Wohlfarth astroid ($b = c = 1$, $\beta = \gamma = 3$).

Ensembles of noninteracting spins ($\alpha = 0$ in flatspin) are subjected to a rising, quasistatic magnetic field aligned with their easy axis, which is held at each field value until equilibrium. Here, equilibrium is defined by the convergence of the mean magnetization over time.

The time intervals were set short enough to avoid a significant probability of multiple flips of one magnet within one interval, and long enough to reach equilibrium within reasonable simulation time. For the low coercivity scenario (1), $\Delta t = 1 \times 10^{-10}$ s, and for the high coercivity scenario (2), $\Delta t = 1 \times 10^{-9}$ s. Other parameters include the attempt frequency, $f_0 = 1 \times 10^9$ Hz and no disorder.

For the low coercivity scenario, the $M(H)$ curve should match the analytical two-state model described by the Brillouin function for spin- $\frac{1}{2}$ systems [33]. The average magnetization of such a system under an applied field $\mu_0 H$ is described by the analytical expression

$$\langle m_x \rangle = \tanh(A\mu_0 H), \quad (13)$$

where $A = M_S V / k_B T$, i.e., $A\mu_0 H$ is the ratio of the Zeeman energy to thermal energy.

For the high coercivity scenario, the energy barriers and the shape of the astroid becomes significant, and the analytical model breaks down. In this case, we compare results with micromagnetic simulations using MuMax3 [21]. The micromagnetic simulations are set up to capture the $M(H)$ curves

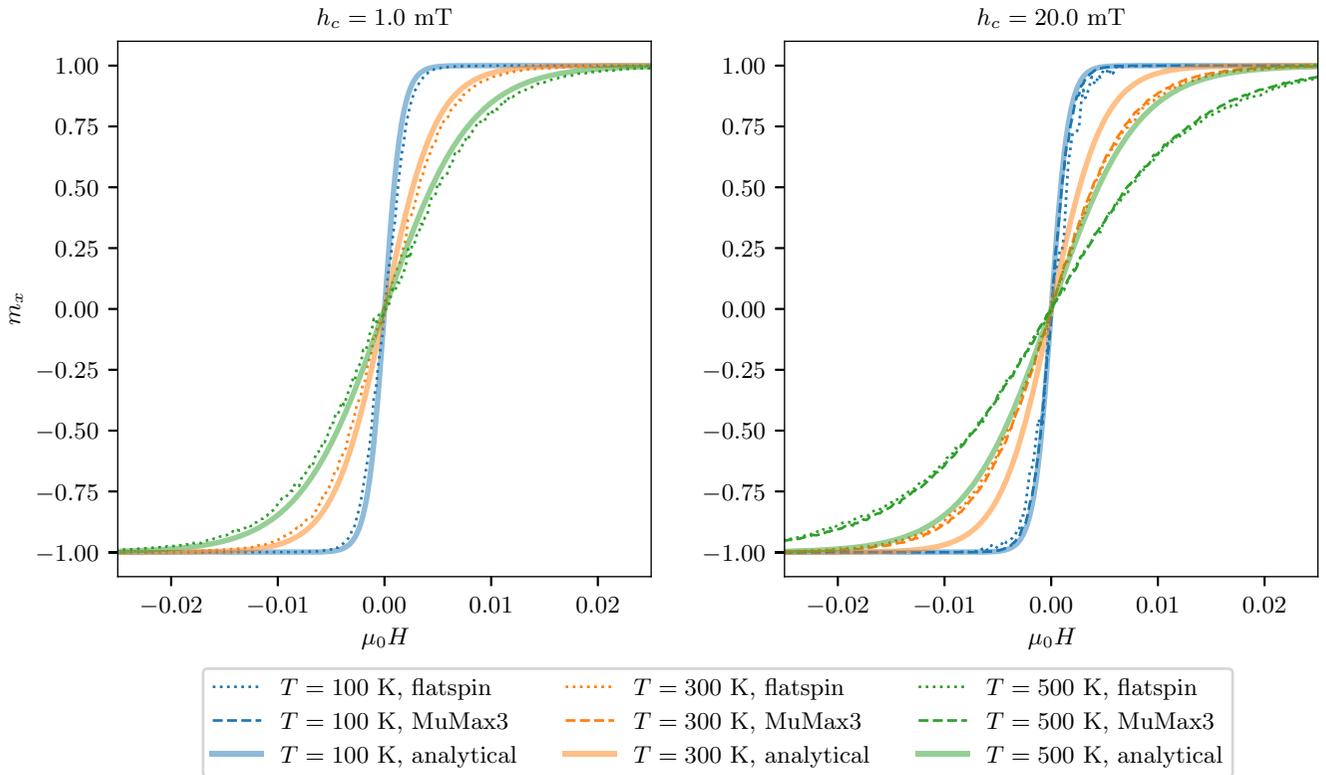


FIG. 6. $M(H)$ curves of ensembles of noninteracting magnets at different temperatures, simulated with different approaches, and for different coercivities. Also indicated is the analytical $\tanh(A\mu_0 H)$ where $A = M_S V/k_B T$. Note that the MuMax3 results are for binned cell magnetization.

of single cell magnets in a rising magnetic field. To isolate the cells in the MuMax3 simulation, the `EnableDemag` property is set to `False` and the exchange stiffness to zero ($A_{\text{ex}} = 0 \text{ J m}^{-1}$). One period of a 500 Hz stepped ramping field is applied, starting at $\mu_0 H_{\text{ext}} = -100 \text{ mT}$ and ending at $\mu_0 H_{\text{ext}} = 100 \text{ mT}$. This field protocol subjects the magnets to a constant field value for $2 \mu\text{s}$ before increasing the field by an additional 0.2 mT , which is sufficient to reach thermal equilibrium. The coercivity of $h_c = 20 \text{ mT}$ was reproduced with a material uniaxial anisotropy and $K_{u1} = h_c \cdot M_S/2 = 8600 \text{ J m}^{-3}$. Other parameters include a world size of 16×16 cells and a cell size of $V = 10 \text{ nm}^3 \times 10 \text{ nm}^3 \times 10 \text{ nm}^3$.

Unlike in flatspin, where spins are binary, the cells in the micromagnetic simulations are allowed to exhibit any magnetization direction. To compare the results from flatspin and micromagnetic simulations, the magnetization of micromagnetic cells is binned into spin states of $s_i = \pm 1$, before they are averaged. Note that this binning only approximates the same average magnetization as a system with significant anisotropy, such as in the high coercivity scenario.

For all simulations, we use the parameters $M_S = 860 \text{ kA m}^{-1}$, and $V = 10 \text{ nm} \times 10 \text{ nm} \times 10 \text{ nm}$.

Figure 6 presents the results of the flatspin simulations, the micromagnetic simulations, as well as the analytical two-state model of Eq. (13). For the low coercivity scenario ($h_c = 1.0 \text{ mT}$), the $M(H)$ curves produced by flatspin agree well with the analytical two-state model. For the high coercivity scenario ($h_c = 20.0 \text{ mT}$), there is a significant deviation between flatspin and the two-state model, but excellent agreement between flatspin and micromagnetic simulations.

For the low coercivity scenario, the agreement between flatspin and the analytical two-state model shows that flatspin's thermal field scales correctly compared to the absolute temperature.

As mentioned, for the high coercivity scenario, the analytical model breaks down, and both flatspin and micromagnetic results deviate significantly from the analytical two-state model. However, there is excellent agreement between the results from flatspin and micromagnetic simulations. Both of these models take the effects of a significant coercivity (significant uniaxial anisotropy) into account. Thus, flatspin is shown to reproduce correct thermal activity also with significant coercivity.

The results presented here are all from systems of noninteracting magnets, where the only influences are the external field and the thermal field. However, since interactions between magnets are mediated by dipolar magnetic fields, the results are also valid for systems of dipolar coupled magnets. The additional dipolar fields can be seen as simple additions to the total field at each magnet, and thus does not alter the validity of the temperature model.

In conclusion, our results show excellent agreement between flatspin and the expected thermal activity for both high and low coercivity scenarios, at equilibrium. For the low coercivity scenario, the influence of the switching astroid is negligible and flatspin matches the analytical two-state model. For the high coercivity scenario, where the energy barriers are significant, flatspin shows excellent agreement with micromagnetic simulations. These results thereby validate the scaling of temperature in the thermal model of flatspin (Sec. II E).

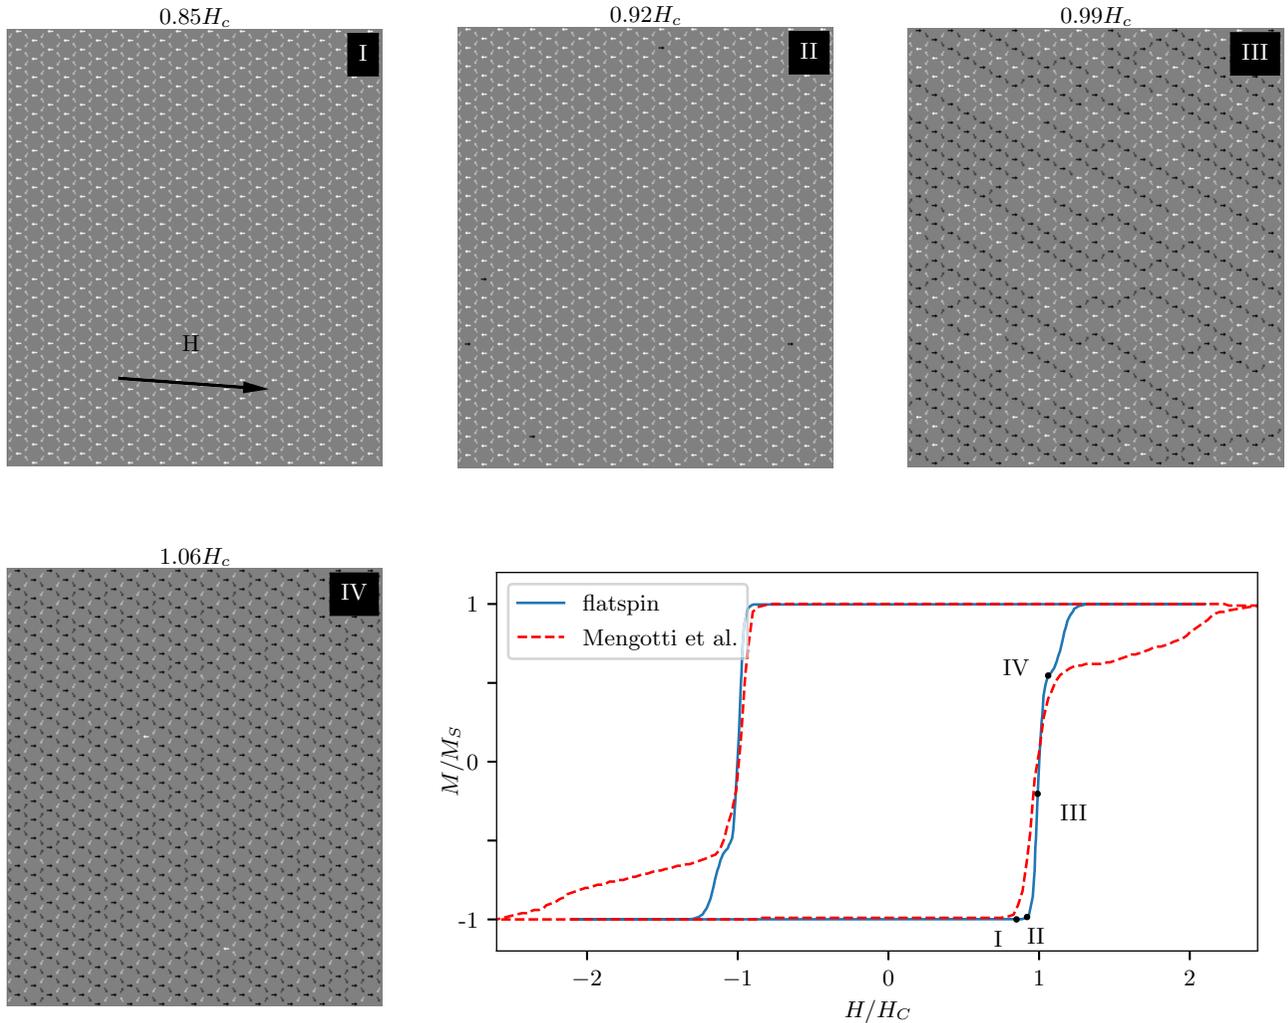


FIG. 7. Left: Snapshots of the evolution of a kagome ASI at selected field values. the images have been cropped to show the middle 80% of the total ASI to improve clarity. Right: Comparison of the hysteresis curve of the simulated ensemble (blue line) against a sketch of the hysteresis curve from the experimental results of Mengotti *et al.* [34] (red-dashed line). The labeled points indicate the points at which the snapshots are sampled from.

V. VALIDATION OF FLATSPIN

To evaluate the suitability of the simulator, flatspin simulations were compared to established experimental results from literature, as well as micromagnetic simulations. In particular, we investigate phenomena such as Dirac strings in kagome ASI [34], the size of crystallite domains in square ASI [10], and superferromagnetism in pinwheel ASI [13]. 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. 4(c)]. Applying a reversal field to a polarized kagome ASI results in the formation of monopole-antimonopole pairs [34]. 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 [34], in which a room temperature kagome ASI undergoes magnetization reversal. We start with an array of 2638 magnets (29×29 hexagons) polarized to the left and apply a reversal field \mathbf{H} to the right with a slight, downward offset of 3.6° . 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 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 time interval $\Delta t = 1 \text{ s}$, and with 5% disorder.

The time evolution snapshots of Fig. 7 demonstrate a strong, qualitative similarity to the results of Mengotti *et al.* [34]. 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 are on the unfavored sublattice, in accordance with both experimental and simulated results from the literature.

Also in Fig. 7, we see the hysteresis of the simulated ensemble (solid line) is similar to that of Mengotti *et al.* [34] (dashed line) in some sections, but differs near the extrema. The hysteresis can be understood in two stages. The first stage, 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 stage in which the Dirac strings have fully covered the array, and the change in magnetization is dominated by switching on the unfavored sublattice. Clearly we see good agreement, within stage one, between our simulated hysteresis and the experimental results. Furthermore, there is a clear transition (characterized by a sharp decrease in gradient) in our hysteresis very close to the transition in the experimental hysteresis. Notably, however, although the transitions occur at similar field values, the change in gradient is less pronounced in our simulated hysteresis. This disparity indicates that, in the second stage, the magnets on the unfavored sublattice flip more easily in our simulation than in the experimental data.

As discussed in Sec. II J, 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 [35], which may contribute to the disparity noted above. Despite this, we observe flatspin accurately reproduces snapshots of the time evolution behavior, while also capturing salient features of the ensemble hysteresis curve.

B. 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.* [10]. One of the main advantages of flatspin over typical alternatives is the scalability 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 spacing is, in part, limited by the number of directly observable magnets.

For a given range of lattice spacings 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. 4(a)] with different lattice spacings, a , ranging from 320 nm to 880 nm.

Square ASIs of size 50×50 were annealed in flatspin with a linearly decreasing temperature, starting at $T = 800$ K and decreasing by 1 K until no magnets were active. Each temperature was simulated over 50 simulation steps. The time interval Δt of each simulation step was chosen so that the

probability of multiple switching events in any single magnet was small, $P_r(k > 1) < 0.001$, or until the total time per temperature value reached 1 min. At high temperatures ($T \gtrsim 720$ K), the requirement $P_r < 0.001$ causes the total time per temperature value to be less than 1 min. This is a trade off between the number of simulation steps allowed and the likelihood of multiple switching events.

The nucleation moment M_{th} was chosen to match the blocking temperature of the 25-nm-thick magnets reported by Zhang *et al.* [36], where moments are stable below ~ 670 K. 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 $10a$ were used. The temperature dependence of the saturation magnetization was accounted for by adjusting M_{th} according to data reported by Zhang *et al.* [36]. The temperature dependence of h_k was scaled by the same factor. Details of the specific time intervals, temperature and temperature dependent parameter values can be found in flatspin's documentation [19].

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 [10,37].

Typical domain structures and correlation results can be seen in Fig. 8. The domains shown in Fig. 8(a) and the correlation curves in Fig. 8(b) are in good agreement with experimental results [10]. A qualitative comparison of the domain sizes and structures in Fig. 8 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 similar trends and values, where an increase in a leads to low correlation, even between nearest neighbors.

The discrepancy for $a = 320$ nm is not completely unexpected, as the point-dipole approximation is known to underestimate nearest-neighbor interaction for magnets placed very close together [35]. In strongly coupled systems, each spin flip results in a greater change in dipole energy, compared to systems that are less coupled. 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 α 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.

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

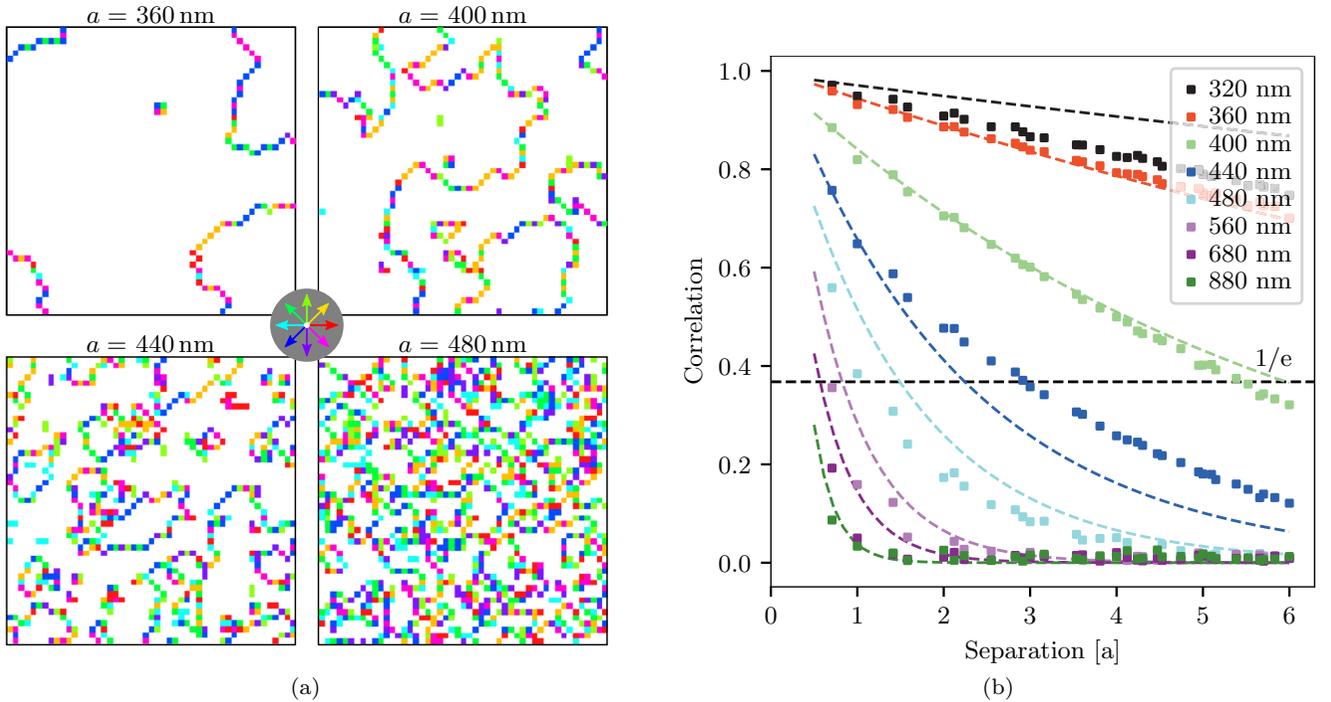


FIG. 8. (a) Maps showing the net magnetization of the ASI vertices, for an annealed 50×50 square ASI with the given lattice spacing a . The white regions have zero net magnetization, and thus correspond to a coherent domain of type I vertices. Colored regions have a nonzero net magnetization, direction 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 spacings, a . Dashed curves show exponential best fits for data from the original paper [10]. Also indicated is a $1/e$ threshold of correlation.

C. 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 [13]. 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 [18]. 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 [13].

Here, we demonstrate that flatspin is able to replicate the experimental reversal processes presented in Li *et al.* [13], where pinwheel “diamond” ASI [Fig. 4(d)] is subject to an external field at different angles. A key result is that the angle θ of the external field controls the nature of the reversal process. When θ is small (equally aligned to both sublattices), reversal happens in a single avalanche, whereas when θ 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 difficult [13].

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. (11). 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 $10a$. Full simulation details are available in the flatspin documentation [19].

Figures 9(a)–9(d) 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. 9(b) and 9(d)] are qualitatively very similar to experimental results [Figs. 9(a) and 9(c)]. In all cases, the ASI undergoes reversal in a single avalanche. Reversal begins at a few 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° as opposed to -6° observed experimentally. Hence, Fig. 9(b) is most similar to Fig. 9(c) and Fig. 9(d) is most similar to Fig. 9(a). It should be noted that the tilted anisotropy axis found experimentally has not yet been explained.

Figures 9(e) and 9(f) show the hysteresis loops and array snapshots when the field is misaligned with the array ($\theta = 30^\circ$). Again, flatspin simulations [Fig. 9(f)] replicate key features observed experimentally [Fig. 9(e)]. 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° .

Li *et al.* [13] report they were unable to replicate the magnetization details using a point-dipole Monte Carlo model.

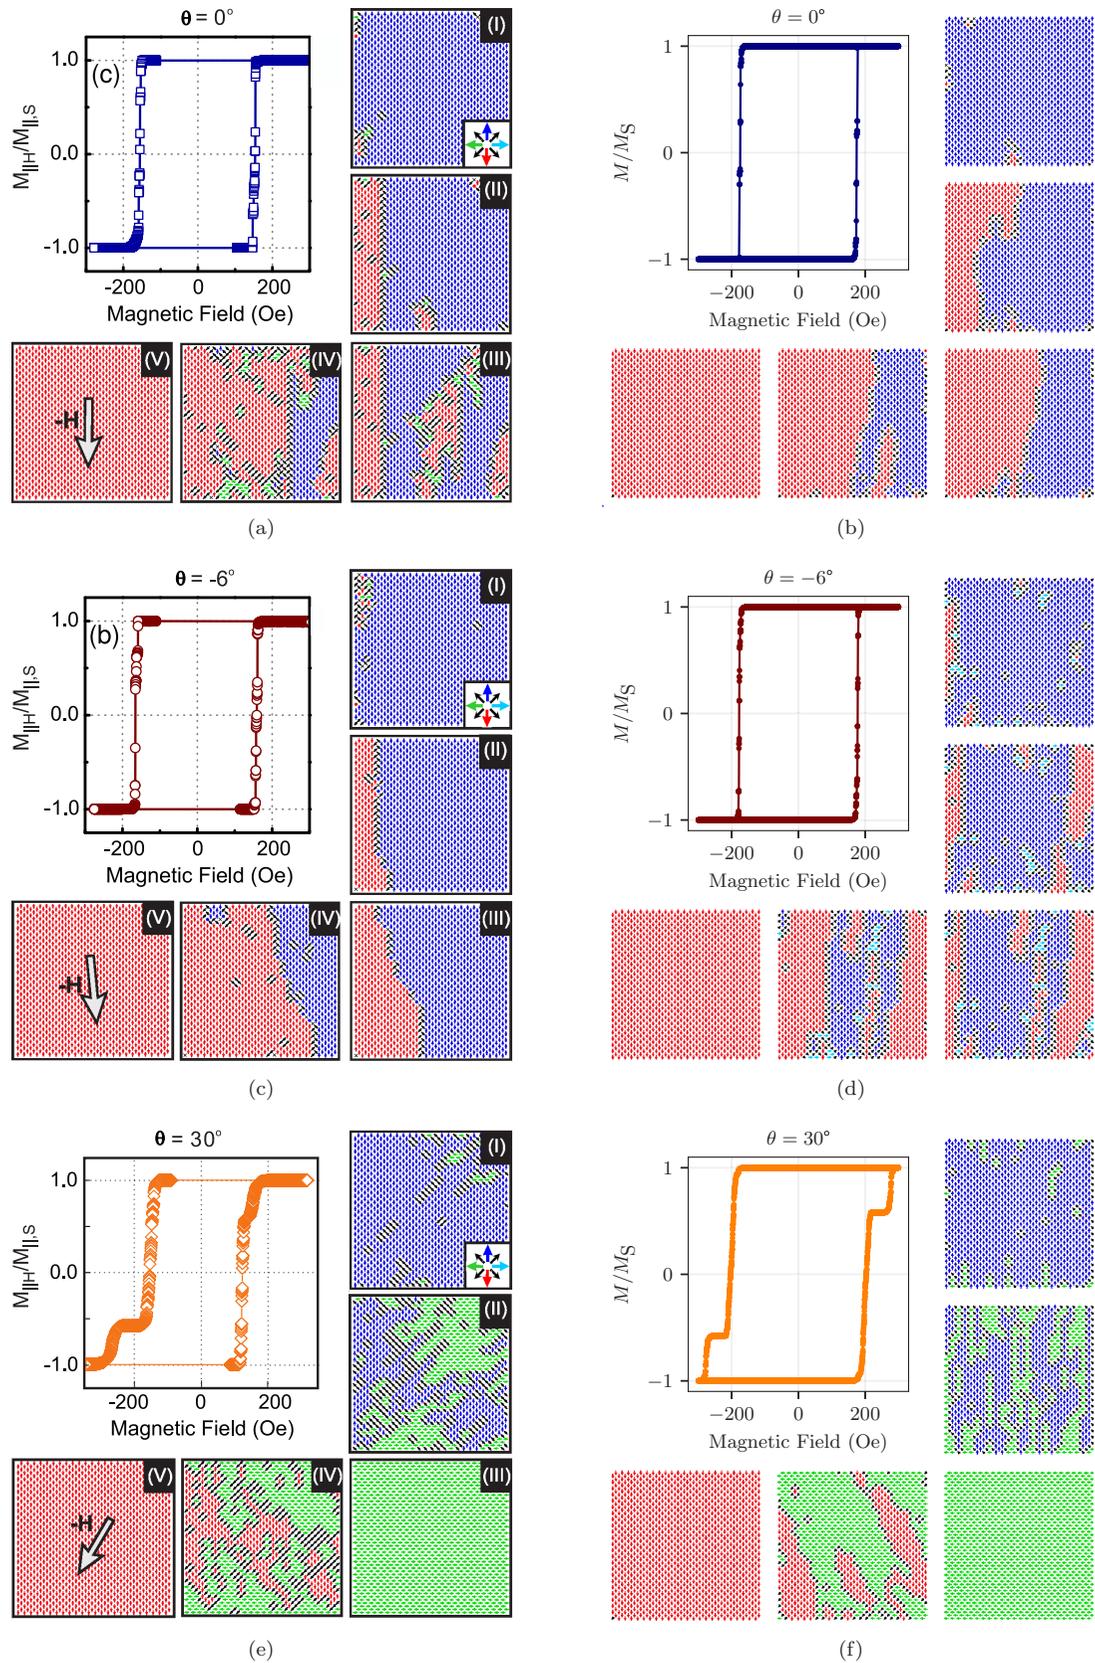


FIG. 9. Hysteresis loop and snapshots of the pinwheel units for various angles θ of the applied field. Figures (a), (c), and (e) show experimental results, adapted from Li *et al.* [13], Copyright ©2018 American Chemical Society, CC-BY-4.0. Figures (b), (d), and (f) show results from flatspin simulation.

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 fields that are not aligned with the magnet's easy axis. Indeed, we find that the generalized Stoner-Wohlfarth model (Sec. II F) is crucial to reproduce the reversal process and magnetization details.

D. Comparison to micromagnetic single-spin switching order

Micromagnetic simulations, e.g., MuMax3 [21,38], are taken as the gold standard and generally agree with experimental results, due to the high simulation fidelity of the micromagnetic model. In this section, we compare how well flatspin agrees with MuMax3 at the level of detail expressed in flatspin.

Here we evaluate the switching strategy outlined in Sec. II H, by comparing the switching orders obtained in flatspin and MuMax3, of a square ASI as it undergoes reversal by an external field. Switching order refers to the sequence in which individual magnets switch their magnetization state. As a similarity measure, Spearman's rank correlation coefficient ρ [39] 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 dominated by the coercivity of each individual magnet, i.e., low coercivity magnets switch first, and high coercivity magnets switch last. 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. In this way, the distribution of magnet coercivities in the two models match.

The system we considered was a 4×4 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 captures switching dynamics, we perform a comparison of the switching orders for all pairs of lattice spacings in both simulators.

Figure 10(a) 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 nonzero y-intercept in the heatmap indicates that, as expected, 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. 10(b) traces the ridge in the heatmap, i.e., the highest ρ , 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 magnets start to interact through dipolar coupling. Below 450 nm, the correlation drops. Since

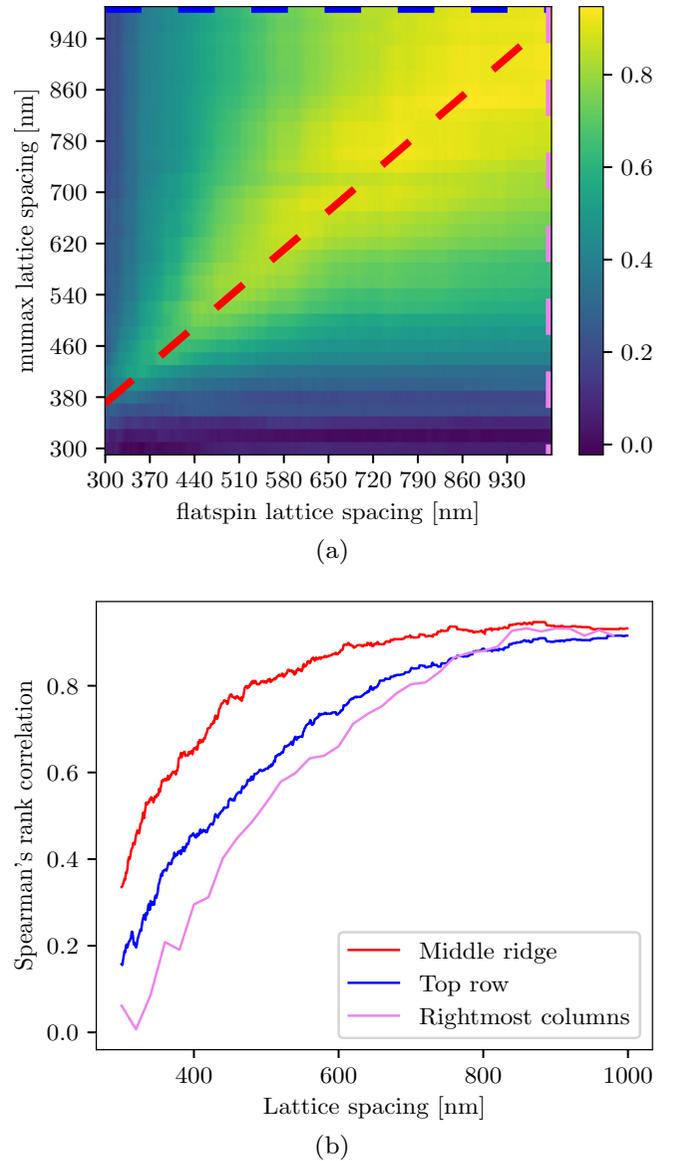


FIG. 10. (a) Spearman's rank correlation coefficients ρ averaged over 32 different square ASIs, evaluated for different lattice spacings in flatspin and MuMax3. The red line shows the approximate maximum ridge line through the heatmap. (b) The red line shows the true maximum ρ for the lattice spacing pairs. The blue and violet lines show projections of the top row and rightmost column of (a), respectively.

flatspin does not account for the micromagnetic state, complete correlation is not expected.

The particular selection of $h_k^{(i)}$ values of flatspin, and the corresponding K_{U1} values of MuMax3, introduces an *inherent* bias in the switching order. One might expect that this quenched disorder dominates the switching order, leading to an inflated correlation between flatspin and MuMax3, regardless of dipole interactions.

The violet line of Fig. 10(b) [plotting the rightmost column from Fig. 10(a)], shows the correlation between MuMax3 and the uncoupled flatspin system (lattice spacing of 1000 nm). If the quenched disorder completely dominated the

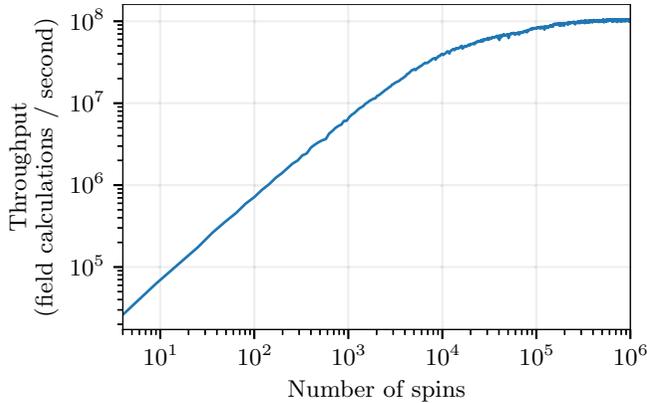


FIG. 11. The throughput (number of field calculations per second) as a function of number of spins. Throughput is averaged over 100 simulations of each size. The test was performed on an NVIDIA Tesla V100 GPU with 32 GB of RAM. Note the logarithmic scale of the axes.

switching order, one would expect this to be a flat line, as increasing the dipole coupling (by reducing lattice spacing) would have no effect. The blue line of Fig. 10(b) [top row from Fig. 10(a)] show the corresponding curve comparing flatspin to the uncoupled MuMax3 system. In both cases, the correlation rapidly declines with lattice spacing, confirming that the switching order is not dominated by the inherent bias for highly coupled systems. Furthermore, the red line clearly shows a stronger agreement when the lattice spacing of MuMax3 and flatspin are both varied proportionally. We conclude that flatspin and MuMax3 capture similar switching dynamics.

VI. 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 .

Figure 11 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 32 GB of RAM.

At around 200 000 spins, the throughput saturates at 10^8 field calculations per second. On our test setup, computing

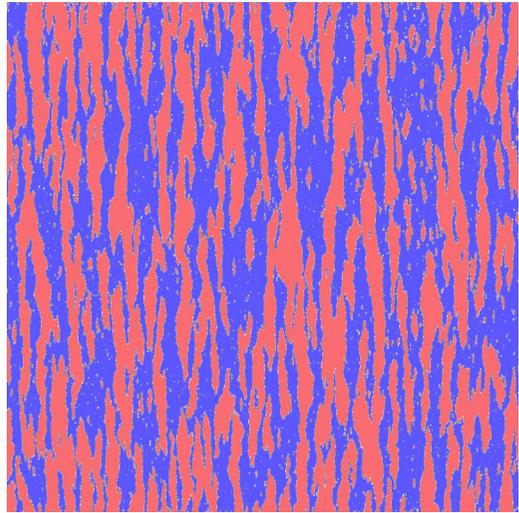


FIG. 12. A snapshot from flatspin simulations of a pinwheel ASI system with more than one million magnets, as it undergoes reversal by an external field. The angle of the external field is $\theta = 0^\circ$.

\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 that results in no spin flips.

The total simulation time depends largely on the particular 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 Sec. VC of 25×25 pinwheel ASI with 1250 magnets took approximately 6 seconds with $K = 2500$, for one reversal.

Figure 12 shows a snapshot from flatspin simulations of a large pinwheel ASI system as it undergoes reversal by an external field. 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 [40].

The ability to simulate such large systems allows a researcher to explore phenomena at much larger scales than can be directly observed experimentally. The imaging window of experimental techniques with single spin resolution, such as magnetic force microscopy (MFM), is typically limited to about $50 \mu\text{m} \times 50 \mu\text{m}$. However, much smaller imaging windows are frequently used due to other practical considerations. For instance, the square ASI investigated by [10] is a typical example, where 5000 magnets are imaged.

VII. 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.

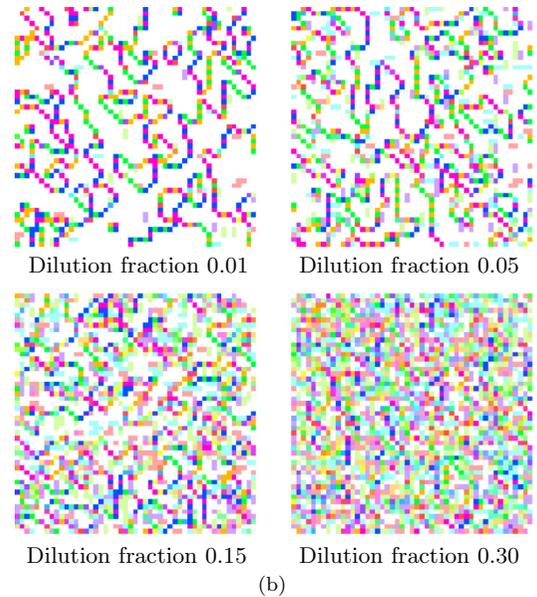
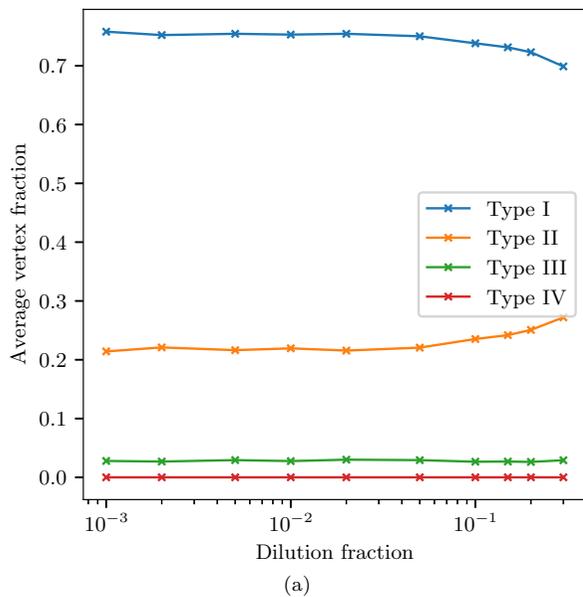


FIG. 13. (a) Average complete vertex fraction, by type, as a function of dilution by removal of spins in a 50×50 square ASI. Only complete vertices are included. The averages are taken over 10 different instances for each dilution fraction. (b) Example end states of square ASI of different dilutions. Each pixel correspond to the net magnetization of 4 spins (vertices), as in Fig. 8(a). Note that white pixels correspond to vertices of no net magnetization (type I vertices *if* the vertices are complete, as there are no type IV vertices).

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 [1,41], 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 (from spin ice I state to the ordered spin ice II state) to disappear [42]. Related work on particle-based ice of the square geometry has revealed an apparent ice rule fragility with respect to similar dilution [43,44]. Here, we use flatspin to investigate the effect of 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 [10]. Reaching this ground-state configuration experimentally 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 50×50 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° , 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 $3a$ is used in the simulations. Larger neighbor distances (up to $50a$) were sampled for the undiluted system and revealed no discrepancy from a neighbor distance of $3a$. 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. 13(a).

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 reduction in domain size, see Fig. 13(b). 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 [10]. 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.

VIII. CONCLUSION

flatspin is a highly effective simulator for ASI systems. At its heart lies a robust magnetic model based on dipole-dipole interactions, with a switching criteria based on a generalized Stoner-Wohlfarth model, and thermal fluctuations based on the Arrhenius-Néel equation. 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 established theory and micromagnetic simulations, and validated

against 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 many ASI simulations to be executed, enabling quick exploration of parameters and novel geometries. The flexibility and performance offered by flatspin open for exciting new possibilities in ASI research.

ACKNOWLEDGMENTS

We thank Bartel Van Waeyenberge for fruitful discussions. This work was funded in part by the Norwegian Research Council by the IKTPLUSS 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 [45]. J.L. is supported by the Research Foundation - Flanders (FWO) through a senior postdoctoral research fellowship (12W7622N).

-
- [1] S. H. Skjærvø, C. H. Marrows, R. L. Stamps, and L. J. Heyderman, Advances in artificial spin ice, *Nat. Rev. Phys.* **2**, 13 (2020).
 - [2] L. J. Heyderman and R. L. Stamps, Artificial ferroic systems: novel functionality from structure, interactions and dynamics, *J. Phys.: Condens. Matter* **25**, 363201 (2013).
 - [3] J. Sklenar, Y. Lao, A. Albrecht, J. D. Watts, C. Nisoli, G.-W. Chern, and P. Schiffer, Field-induced phase coexistence in an artificial spin ice, *Nat. Phys.* **15**, 191 (2019).
 - [4] 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, K. Kiefer, S. Gerischer, D. Slobinsky, and R. S. Perry, Dirac strings and magnetic monopoles in the spin ice $\text{Dy}_2\text{Ti}_2\text{O}_7$, *Science* **326**, 411 (2009).
 - [5] D. Levis, L. F. Cugliandolo, L. Foini, and M. Tarzia, Thermal Phase Transitions in Artificial Spin Ice, *Phys. Rev. Lett.* **110**, 207206 (2013).
 - [6] L. Anghinolfi, H. Luetkens, J. Perron, M. G. Flokstra, O. Sendetskyi, A. Suter, T. Prokscha, P. M. Derlet, S. L. Lee, and L. J. Heyderman, Thermodynamic phase transitions in a frustrated magnetic metamaterial, *Nat. Commun.* **6**, 8278 (2015).
 - [7] J. H. Jensen, E. Folven, and G. Tufte, Computation in artificial spin ice, in *The 2018 Conference on Artificial Life* (MIT Press, Cambridge, MA, 2018), pp. 15–22
 - [8] J. H. Jensen and G. Tufte, Reservoir computing in artificial spin ice, in *The 2020 Conference on Artificial Life* (MIT Press, Cambridge, MA, 2020), pp. 376–383.
 - [9] Z. Budrikis, Athermal dynamics of artificial spin ice: disorder, edge and field protocol effects, Ph.D. thesis, The University of Western Australia, 2012.
 - [10] S. Zhang, I. Gilbert, C. Nisoli, G. W. Chern, M. J. Erickson, L. O'Brien, C. Leighton, P. E. Lammert, V. H. Crespi, and P. Schiffer, Crystallites of magnetic charges in artificial spin ice, *Nature (London)* **500**, 553 (2013).
 - [11] X. Ke, J. Li, C. Nisoli, P. E. Lammert, W. McConville, R. F. Wang, V. H. Crespi, and P. Schiffer, Energy Minimization and ac Demagnetization in a Nanomagnet Array, *Phys. Rev. Lett.* **101**, 037205 (2008).
 - [12] Y. Bar-Yam, *Dynamics of Complex Systems* (Addison-Wesley, Reading, MA, 1997).
 - [13] Y. Li, G. W. Paterson, G. M. Macaulay, F. S. Nascimento, C. Ferguson, S. A. Morley, M. C. Rosamond, E. H. Linfield, D. A. MacLaren, R. Macêdo *et al.*, Superferromagnetism and domain-wall topologies in artificial “pinwheel” spin ice, *ACS Nano* **13**, 2213 (2019).
 - [14] J. Frenkel and J. Dorfman, Spontaneous and induced magnetisation in ferromagnetic bodies, *Nature (London)* **126**, 274 (1930).
 - [15] C. Kittel, Theory of the structure of ferromagnetic domains in films and small particles, *Phys. Rev.* **70**, 965 (1946).

- [16] R. F. Wang, C. Nisoli, R. S. Freitas, J. Li, W. McConville, B. J. Cooley, M. S. Lund, N. Samarth, C. Leighton, V. H. Crespi, and P. Schiffer, Artificial ‘spin ice’ in a geometrically frustrated lattice of nanoscale ferromagnetic islands, *Nature (London)* **439**, 303 (2006).
- [17] C. Fan and F. Y. Wu, Ising model with second-neighbor interaction. I. Some exact results and an approximate solution, *Phys. Rev.* **179**, 560 (1969).
- [18] R. Macêdo, G. M. Macauley, F. S. Nascimento, and R. L. Stamps, Apparent ferromagnetism in the pinwheel artificial spin ice, *Phys. Rev. B* **98**, 014437 (2018).
- [19] flatspin contributors, flatspin website, <https://flatspin.gitlab.io> (2020).
- [20] L. Néel, Théorie du traînage magnétique des ferromagnétiques en grains fins avec applications aux terres cuites, *Ann. Geophys.* **5**, 99 (1949).
- [21] A. Vansteenkiste, J. Leliaert, M. Dvornik, M. Helsen, F. Garcia-Sanchez, and B. Van Waeyenberge, The design and verification of MuMax3, *AIP Adv.* **4**, 107133 (2014).
- [22] C. Tannous and J. Gieraltowski, The stoner–wöhlfarth model of ferromagnetism, *Eur. J. Phys.* **29**, 475 (2008).
- [23] R. Kikuchi, On the minimum of magnetization reversal time, *J. Appl. Phys.* **27**, 1352 (1956).
- [24] P. R. Gillette and K. Oshima, Magnetization reversal by rotation, *J. Appl. Phys.* **29**, 529 (1958).
- [25] M. Tanaka, E. Saitoh, H. Miyajima, T. Yamaoka, and Y. Iye, Magnetic interactions in a ferromagnetic honeycomb nanoscale network, *Phys. Rev. B* **73**, 052411 (2006).
- [26] Y. Qi, T. Brintlinger, and J. Cumings, Direct observation of the ice rule in an artificial kagome spin ice, *Phys. Rev. B* **77**, 094418 (2008).
- [27] M. Saccone, A. Scholl, S. Velten, S. Dhuey, K. Hofhuis, C. Wuth, Y.-L. Huang, Z. Chen, R. V. Chopdekar, and A. Farhan, Towards artificial Ising spin glasses: Thermal ordering in randomized arrays of Ising-type nanomagnets, *Phys. Rev. B* **99**, 224403 (2019).
- [28] A. Libál, C. Reichhardt, and C. J. Olson Reichhardt, Realizing Colloidal Artificial Ice on Arrays of Optical Traps, *Phys. Rev. Lett.* **97**, 228302 (2006).
- [29] A. Ortiz-Ambriz and P. Tierno, Engineering of frustration in colloidal artificial ices realized on microfeatured grooved lattices, *Nat. Commun.* **7**, 10575 (2016).
- [30] P. Mellado, A. Concha, and L. Mahadevan, Macroscopic Magnetic Frustration, *Phys. Rev. Lett.* **109**, 257203 (2012).
- [31] F. Ma, C. Reichhardt, W. Gan, C. J. Olson Reichhardt, and W. S. Lew, Emergent geometric frustration of artificial magnetic skyrmion crystals, *Phys. Rev. B* **94**, 144405 (2016).
- [32] A. Libál, C. J. Olson Reichhardt, and C. Reichhardt, Creating Artificial Ice States Using Vortices in Nanostructured Superconductors, *Phys. Rev. Lett.* **102**, 237004 (2009).
- [33] C. Kittel, *Introduction to Solid State Physics*, 8th ed. (John Wiley, Hoboken, NJ, 2004), pp. 303–304.
- [34] E. Mengotti, L. J. Heyderman, A. F. Rodríguez, F. Nolting, R. V. Hügli, and H.-B. Braun, Real-space observation of emergent magnetic monopoles and associated dirac strings in artificial kagome spin ice, *Nat. Phys.* **7**, 68 (2011).
- [35] N. Rougemaille, F. Montaigne, B. Canals, A. Duluard, D. Lacour, M. Hehn, R. Belkhou, O. Fruchart, S. El Moussaoui, A. Bendounan, and F. Maccherozzi, Artificial Kagome Arrays of Nanomagnets: A Frozen Dipolar Spin Ice, *Phys. Rev. Lett.* **106**, 057209 (2011).
- [36] X. Zhang, Y. Lao, J. Sklenar, N. S. Bingham, J. T. Batley, J. D. Watts, C. Nisoli, C. Leighton, and P. Schiffer, Understanding thermal annealing of artificial spin ice, *APL Mater.* **7**, 111112 (2019).
- [37] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Clarendon Press, Oxford, 1999).
- [38] J. Leliaert, M. Dvornik, J. Mulkers, J. De Clercq, M. V. Milošević, and B. Van Waeyenberge, Fast micromagnetic simulations on GPU - Recent advances made with mumax3, *J. Phys. D* **51**, 123002 (2018).
- [39] J. D. Gibbons and S. Chakraborti, *Nonparametric Statistical Inference: Revised and Expanded* (CRC Press, Boca Raton, FL, 2014).
- [40] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevB.106.064408> for a video of the reversal of 708x708 pinwheel ASI.
- [41] L. A. Mól, R. L. Silva, R. C. Silva, A. R. Pereira, W. A. Moura-Melo, B. V. Costa, and M. Carlo, Magnetic monopole and string excitations in two-dimensional spin ice, *J. Appl. Phys.* **106**, 063913 (2009).
- [42] P. Andriushchenko, Influence of cutoff dipole interaction radius and dilution on phase transition in kagome artificial spin ice, *J. Magn. Magn. Mater.* **476**, 284 (2019).
- [43] A. Libál, D. Y. Lee, A. Ortiz-Ambriz, C. Reichhardt, C. J. O. Reichhardt, P. Tierno, and C. Nisoli, Ice rule fragility via topological charge transfer in artificial colloidal ice, *Nat. Commun.* **9**, 4146 (2018).
- [44] C. Nisoli, Unexpected Phenomenology in Particle-Based Ice Absent in Magnetic Spin Ice, *Phys. Rev. Lett.* **120**, 167205 (2018).
- [45] M. Sjölander, M. Jahre, G. Tufte, and N. Reissmann, EPIC: An energy-efficient, high-performance GPGPU computing research infrastructure, [arXiv:1912.05848](https://arxiv.org/abs/1912.05848).