ABSTRACT

Title of Dissertation:	THERMAL ACTIVATION, LONG-RANGE ORDERING, AND TOPOLOGICAL FRUSTRATION OF ARTIFICIAL SPIN ICE
	Jasper Altman Drisko Doctor of Philosophy, 2016
Dissertation directed by:	Professor John Cumings Department of Materials Science and Engineering

Frustrated systems, typically characterized by competing interactions that cannot all be simultaneously satisfied, are ubiquitous in nature and display many rich phenomena and novel physics. Artificial spin ices (ASIs), arrays of lithographically patterned Ising-like single-domain magnetic nanostructures, are highly tunable systems that have proven to be a novel method for studying the effects of frustration and associated properties. The strength and nature of the frustrated interactions between individual magnets are readily tuned by design and the exact microstate of the system can be determined by a variety of characterization techniques. Recently, thermal activation of ASI systems has been demonstrated, introducing the spontaneous reversal of individual magnets and allowing for new explorations of novel phase transitions and phenomena using these systems. In this work, we introduce a new, robust material with favorable magnetic properties for studying thermally active ASI and use it to

investigate a variety of ASI geometries. We reproduce previously reported perfect ground-state ordering in the square geometry and present studies of the kagome lattice showing the highest yet degree of ordering observed in this fully frustrated system. We consider theoretical predictions of long-range order in ASI and use both our experimental studies and kinetic Monte Carlo simulations to evaluate these predictions. Next, we introduce controlled topological defects into our square ASI samples and observe a new, extended frustration effect of the system. When we introduce a dislocation into the lattice, we still see large domains of ground-state order, but, in every sample, a domain wall containing higher energy spin arrangements originates from the dislocation, resolving a discontinuity in the ground-state order parameter. Locally, the magnets are unfrustrated, but frustration of the lattice persists due to its topology. We demonstrate the first direct imaging of spin configurations resulting from topological frustration in any system and make predictions on how dislocations could affect properties in numerous materials systems.

THERMAL ACTIVATION, LONG-RANGE ORDERING, AND TOPOLOGICAL FRUSTRATION OF ARTIFICIAL SPIN ICE

by

Jasper Altman Drisko

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2016

Advisory Committee:
Professor John Cumings, Chair
Professor Christopher Lobb
Professor Richard Greene
Professor Frederick Wellstood
Professor Ichiro Takeuchi

Acknowledgements

First and foremost, I would like to thank my advisor, John Cumings, for his constant support, encouragement, and belief in me during my time at Maryland. He always had valuable insight on my research and pushed me to become the best scientist I could be. He was an encyclopedia of knowledge about both experimental techniques and theoretical ideas. He always respected me, listened to me, and supported what I wanted to do. I am grateful for everything I was able to learn from him.

I would also like to thank all of the great teachers I have been fortunate enough to have throughout my life who inspired me and contributed to my love of learning. There are too many to list, but I am especially thankful for the influence that Mark Powell, Sue Davenport, Biff Maier, Dr. Michael Fiveash, Todd Kresser, Nick Gould, Dr. Bryan Partridge, Dr. D. Barnes Boffey, Dr. Alex Zettl, Dr. Paulo Bedaque, and Dr. Ian Appelbaum had in my education.

I would like to thank my committee members, Dr. Rick Greene, Dr. Chris Lobb, Dr. Fred Wellstood, and Dr. Ichiro Takeuchi, for their support and evaluation of my work. I am grateful to Ichiro for access to his sputter chamber which I used to anneal my thin films, without which this dissertation would not have been possible. I would also like to thank Rick for the experience of working with him as a TA, which I greatly enjoyed.

I want to thank my undergraduate research mentors, Dr. Shane Cybart and Dr. Bob Dynes, for introducing me to experimental research and arming me a with a wealth of knowledge and skills for graduate school and beyond. I want to thank Dr. Stephen Wu for valuable advice about navigating graduate school.

I am extremely grateful to the other members of Cumings' group with whom I had the pleasure of collaborating and working with. Stephen Daunheimer worked on the artificial spin ice project before me and gave me priceless information, training, and tips about the project. I am grateful to him for teaching me the ropes and answering all my questions when I was starting this research. I also want to thank Thomas Marsh, a high school student who worked closely with me for a summer and greatly contributed to the topological frustration project. I am grateful to Hanna Nilsson for her friendship, moral support, listening skills, motivation, and useful discussions over the years. I also want to acknowledge valuable discussions, advice, and help from all the other members of our group: Kai He, Merijntje Bronsgeest, Vladimir Oleshko, Kamal Baloch, Khim Karki, John Bavier, Norvik Voskanian, Paris Alexander, Jeremy Ticey, Zoey Warecki, Eric Epstein, Konnor Kim, Ben Shapiro, and Raj Bajwa. I want to thank Kale Johnson, Sean Fackler, and Alex Kozen for their friendship and enlightening discussions.

I am very grateful to all of the staff at Maryland I was able to interact with and who taught and helped me immensely along the way. I want to thank the AIMLab staff, especially Dr. Phil Piccoli and Dr. Wen-An Chiou, the FABLab staff, especially Tom Loughran and Jon Hummel, the CNAM staff, Brian Straughn and Doug Bensen, and the MSE business office staff, especially Ginette Villeneuve.

I want to thank my bandmates in The Vico Cycle, B Side Shuffle, and Staycation, and everyone I was fortunate enough to play music with during my time at Maryland. Music served as a perfect distraction and balance to the long days in the lab.

Finally, I want to thank my mom, dad, and sister Arielle, for unconditionally supporting me throughout my education and my life. I love you.

Table of Contents

List of Figures

Chapter 1 Introduction and background

Artificial spin ice (ASI) is a highly tunable model system that can be used to gain unique insights into frustrated systems. In this chapter, I will introduce a variety of topics relevant to ASI. I will begin with a definition of frustration and describe the first crystalline system discovered to be frustrated, water ice. I then describe the quantum origins and different types of conventional magnetism including frustrated magnetism. After a brief survey of different classes of magnetically frustrated materials, an introduction to the pyrochlore spin ices will be given. Finally, I provide a thorough overview of ASI, the main focus of this dissertation, highlighting seminal initial work and describing some recent results on the thermal activation of ASI systems.

1.1 Frustration and water ice

Frustrated systems are characterized by competing interactions that cannot all be simultaneously satisfied. They are ubiquitous in nature and display many novel phenomena not found elsewhere in physics. Many physical systems tend to find unique, highly-ordered, low energy, and low entropy configurations as the temperature of the system is lowered towards 0 K. The Third Law of Thermodynamics, originally known as the Nernst heat theorem [1], states that as the temperature of a system tends towards 0 K, the entropy of the system, a measure of its disorder, should also go to 0. Frustrated systems, however, do not obey this law. A high number of degenerate or nearly degenerate ground states is

characteristic of a frustrated system. There is built-in disorder that persists down to the lowest temperatures, resulting in a residual entropy.

The first crystalline system known to exhibit frustration is water ice, which was described by Pauling in 1935 [2] to resolve a discrepancy in different methods of calculating the entropy of ice [3,4]. Entropy can be expressed as

$$S(T) = \int_0^T \frac{C_p(T')dT'}{T'} \tag{1.1}$$

where S is the entropy, T is the temperature, and C_p is the heat capacity at constant pressure. By measuring the heat capacity of a system, one can use Equation 1.1 to find the entropy. Giauque and Ashley did exactly this [3], using heat capacity measurements of ice at low temperatures to calculate the entropy of the system [5,6]. A separate method to find the entropy of a system is to use spectroscopic band data of a gas and knowledge of the corresponding molecular structure to calculate the total entropy [7–9]. Gordon calculated entropy in this manner for water vapor and found a value that did not match the number found from heat capacity measurements [4].

Pauling showed that this discrepancy in the entropy is actually due to the degenerate ground state of the ice system. As temperature is lowered towards 0 K, the oxygen atoms form a tetrahedral lattice where each oxygen is bonded to four other oxygens around it (see Figure 1.1) [10]. The hydrogen atoms sit on the oxygen-oxygen bonds, but do not sit in the middle; each hydrogen atom is closer to one oxygen atom or the other. The closer hydrogen atoms are 0.95 Å from the oxygen and are covalently bonded, corresponding to an H_2O molecule, while the further atoms are 1.81 Å apart and only hydrogen bonded [11]. Pauling was aware of both the structure of oxygen atoms and the oxygen-hydrogen bond lengths from x-ray measurements. He assumed that the lowest energy configuration, locally, has

two hydrogen atoms close to each oxygen and two further away, a structure proposed by Bernal and Fowler and known as the "two-in, two-out" ice rules [11]. Pauling showed that there are 6 different ways to satisfy the ice rules out of $2^4 = 16$ total possible configurations for the 4 hydrogen atoms around each oxygen. He then solved what would now be an undergraduate statistical mechanics problem to calculate the residual entropy in ice due to the 6 out of 16 possible lowest energy configurations.

For N oxygen atoms there are 2^{2N} possible states, since each oxygen shares a bond with 4 other oxygens, and $(6/16)^N$ of these states obey the ice rules. Therefore the total number of possible ground-state configurations is

$$\Omega_0 = \left[2^{2N}\left(\frac{6}{16}\right)^N\right] \tag{1.2}$$

and the residual entropy is

$$S_0 = k_B \cdot ln(\Omega_0) = k_B \cdot ln\left(\left(\frac{3}{2}\right)^N\right) = Nk_B \cdot ln\left(\frac{3}{2}\right) \tag{1.3}$$

For N atoms in a mole, Nk_B becomes the gas constant R = 1.987 cal $K^{-1}mol^{-1}$ and then S_0 = 0.805 cal $K^{-1}mol^{-1}$, which is very close to the difference in entropy of 0.870 cal K^{-1} mol^{-1} between the values reported in Refs. [3] and [4]. Further, careful heat capacity measurements reduced the measured value of the excess entropy to 0.820 cal K^{-1} mol^{-1} [6] and more rigorous calculations resulted in a theoretical value of 0.814 cal K^{-1} mol^{-1} [12]. It is interesting to note that the same amount of residual entropy per atom is present in the magnetic pyrochlore spin ices which I discuss in detail in Section 1.3 [13].

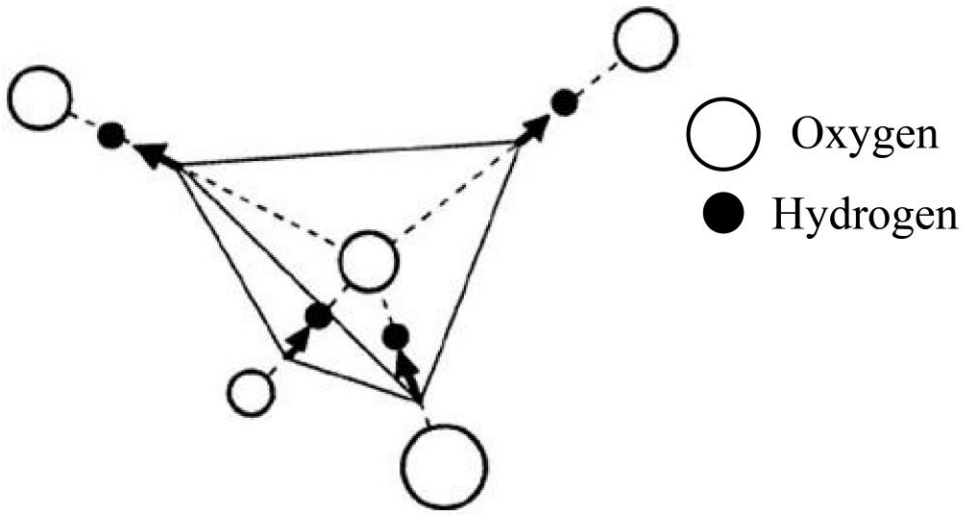

Figure 1.1 **Tetrahedral structure of ice**. As ice is cooled towards 0 K, the oxygen atoms (open circles) form a tetrahedral structure where each oxygen is bonded to four others around it. The hydrogen atoms (filled circles) sit on the oxygen-oxygen bonds, but don't sit in the middle. Each hydrogen proton is closer to one oxygen atom or the other. This schematic shows a lowest energy "ice rule" state, where two of the hydrogens are close to the center oxygen and two are further away. There are 6 different ways to satisfy this 2-in, 2-out local ground state ordering, leading to geometric frustration. Figure adapted from Ref. [14].

A great number of other frustrated systems have since been discovered and the frustrated interactions can arise from a variety of sources. Many of the interactions are magnetic in nature as will be discussed at length in this dissertation. Some other examples of non-magnetic systems with residual entropy are molecular crystals, such as crystalline CO, where disorder comes from varying orientations of the molecules, and in quenched alloys, such as β-brass (CuZn), where rapid cooling results in disordered positions of atoms throughout the lattice [15].

1.2 Magnetism

1.2.1 Conventional magnetism

Electric and magnetic fields are well described by Maxwell's equations and their intrinsic relationship is an essential topic of classical physics. In materials, however, the fundamental, atomic origin of magnetism is much more complex and generally quantum mechanical in nature. In the most basic sense, it arises from the quantized angular momentum, either orbital or intrinsic (spin), of electrons around atoms and whether the atoms have completely filled or partially filled outer electronic shells. Different types of magnetism in materials and diverse magnetic behaviors originate from a variety of (sometimes competing) factors. In this section I provide an overview of these factors and a survey of the resultant bulk properties. The physics of small ferromagnetic islands with strong shape anisotropy will also be discussed, as these will be especially important in later chapters.

The magnetic field $\boldsymbol{B}$ inside a material is related to the external magnetic induction $\boldsymbol{H}$ and the magnetization of the material (the magnetic moment per unit volume) $\boldsymbol{M}$ by

$$\boldsymbol{B} = \mu_0(\boldsymbol{H} + \boldsymbol{M}) \tag{1.4}$$

$\boldsymbol{M}$ is related to $\boldsymbol{H}$ as

$$\boldsymbol{M} = \chi \boldsymbol{H} \tag{1.5}$$

where χ is the magnetic susceptibility, a property which can vary greatly between materials and is also a function of temperature [16]. $\chi < 0$ is defined as diamagnetism, where the magnetization of a material opposes an applied field. This can be understood as a microscopic application of Lenz's law, in which a changing applied magnetic field induces

a current in a loop of wire such that the induced field from that current will oppose the change in applied field. In the classical picture, when a field is applied to a material, electrons will begin to orbit due to Larmor precession [17]. The induced electron orbits will generate a magnetic field that tends to oppose the applied field. The same result can be similarly obtained from quantum mechanics. Interestingly, all materials have a weak diamagnetic contribution, though if another type of magnetism is present the diamagnetic part will generally be negligible. Many purely diamagnetic materials typically have susceptibilities around $\chi \approx -10^{-5}$ [18]. This is very weak, though the application of a sufficiently strong magnetic field (>10 T) can induce enough magnetism to allow many ordinary, seemingly non-magnetic objects to levitate in the applied field, such as a frog or a hazelnut as was demonstrated by Berry and Geim [18]. Bulk superconductors display perfect diamagnetism, the Meissner effect [19], though the microscopic origin is different than for the atoms discussed above. Persistent screening currents in the superconductor generate fields that cancel (either perfectly in type-I superconductors or partially in type-II) the applied magnetic field, though the application of a strong enough external field will destroy the superconductivity in either type [20]. A schematic of the Meissner effect is shown in Figure 1.2.

The net magnetic moment of an atom is determined by the sum of the orbital and spin angular momentum [16]. If the atom has full electron shells, it will have no net moment and can only display diamagnetism. If instead, a material has atoms with open shells, there will be a net moment that will tend to align with an applied magnetic field since that will lower the energy of the system. This will result in $\chi > 0$ and is defined to be paramagnetism.

The temperature dependence of χ for a simple paramagnet is given by the Curie law [17]

$$\chi = \frac{C}{T} \tag{1.6}$$

In metals, where conduction electrons are delocalized from the atoms and energy levels are described by an electronic band structure, the picture is slightly different. Electrons of opposite spin respond differently to the applied field, which shifts the dispersion relations in opposite directions for each spin orientation. Spins aligned antiparallel to the field have more energy, so the relative number of parallel spins will increase in order to keep the filled electron states at the Fermi level [17]. The preferential population of one of the spin species results in a net magnetization in the material. This is known as Pauli paramagnetism and is shown schematically in Figure 1.3.

Both diamagnetism and paramagnetism are only present during the application of an external field, but atoms can also have a spontaneous magnetic moment in the absence of an applied field. Atoms with a spontaneous moment can also interact with each other through the overlap of their wavefunctions. Different materials favor parallel or antiparallel alignment of atomic moments through the exchange interaction. Exchange arises through a combination of the Pauli exclusion principle and the Coulombic repulsion between electrons. For conduction electrons *i* and *j* of the same spin in a metal, the probability electron *i* will be found in volume $d\boldsymbol{r}_i$ and electron *j* found in volume $d\boldsymbol{r}_j$ is [16]

$$\left|\psi_{ij}\right|^2 d\boldsymbol{r_i} d\boldsymbol{r_j} = \frac{1}{V^2}\left[1 - \cos\left(\left(\boldsymbol{k}_i - \boldsymbol{k}_j\right)\cdot\left(\boldsymbol{r_i} - \boldsymbol{r_j}\right)\right)\right] d\boldsymbol{r}_i d\boldsymbol{r}_j \tag{1.7}$$

This clearly shows that an electron of a given spin orientation cannot be in the same location as a second electron of the same spin as the probability for this to occur is 0.

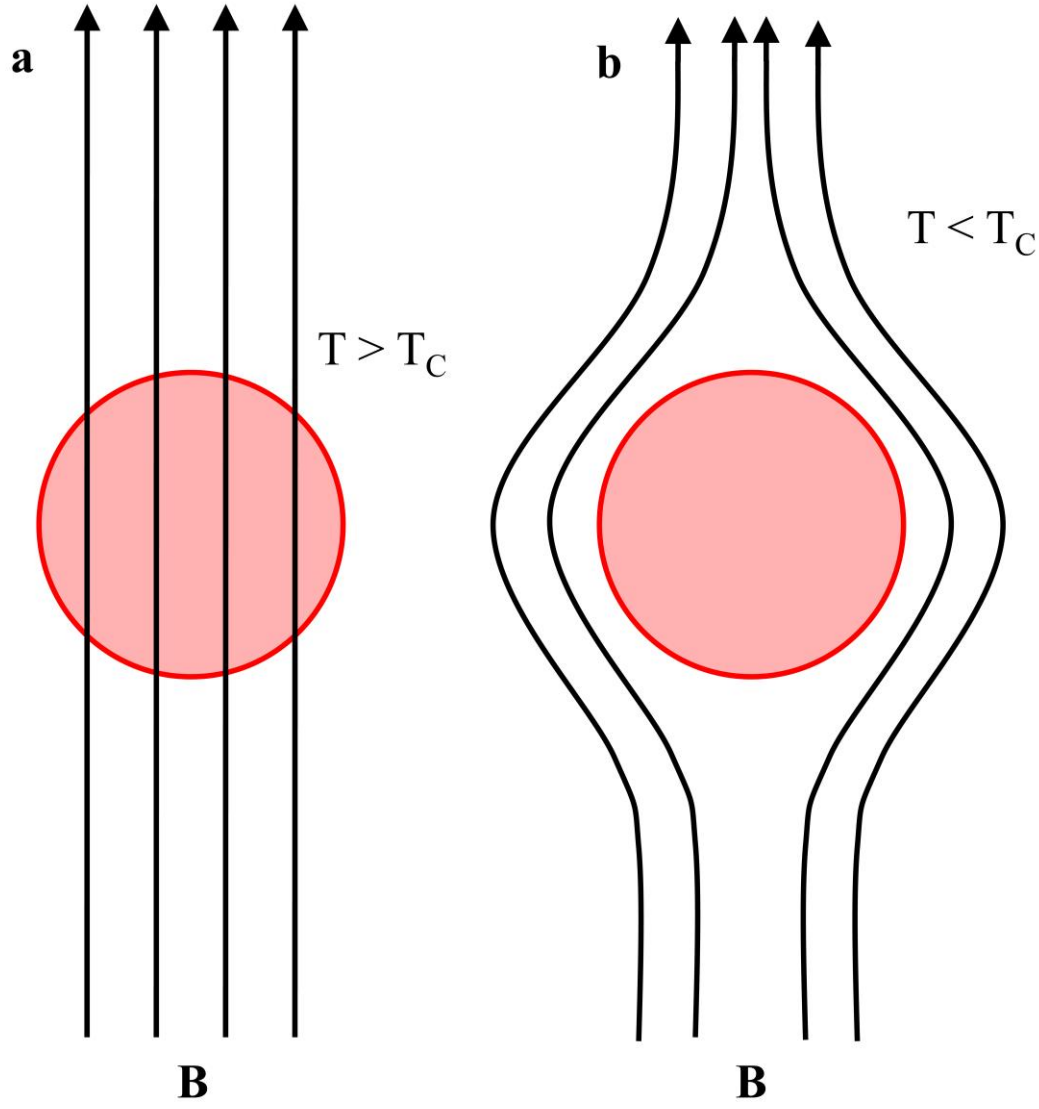

Figure 1.2 **The Meissner effect**. As a superconductor (red circle) is cooled through its transition temperature (T_C), magnetic field lines are expelled from the material. The net field inside the superconductor becomes zero and it behaves like a perfect diamagnet. **a**. Above T_C. **b**. Below T_C.

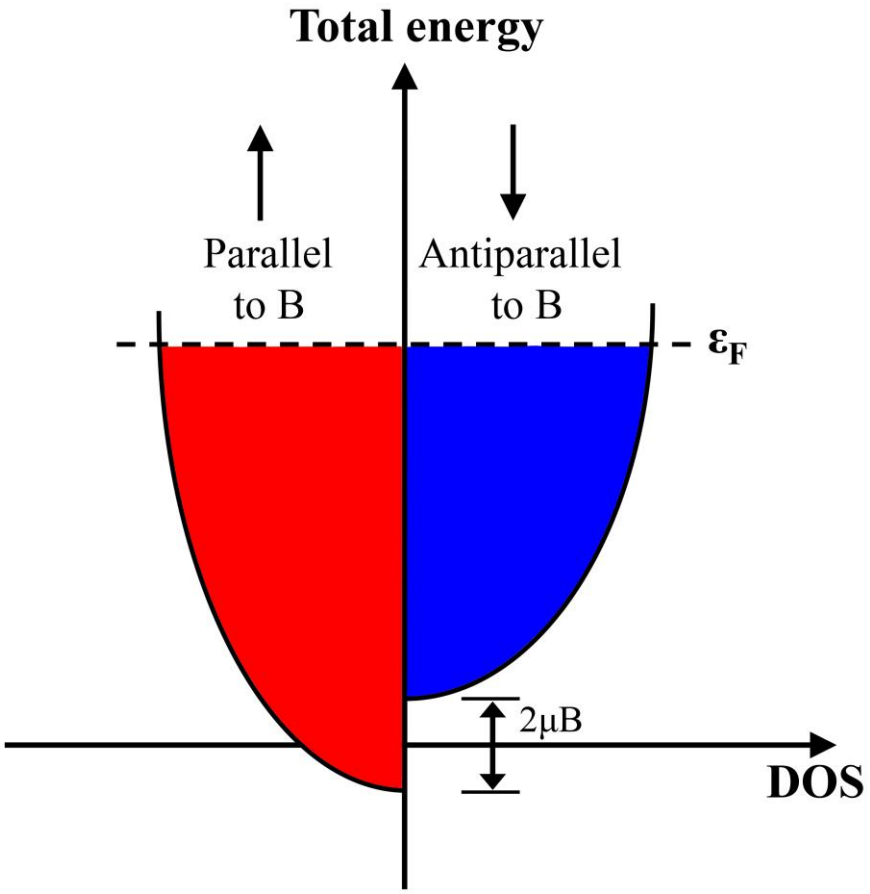

Figure 1.3 **Pauli paramagnetism**. An applied magnetic field induces a higher population of one spin species resulting in a net magnetic moment. Here, the electrons antiparallel to the field cost more energy, so the parallel spins increase in number while the antiparallel decrease so that the total energies of the two populations are equal to the Fermi level, ε_F. μ is the magnetic moment of the electrons.

The exchange interaction is given as [17]

$$U = -2J\boldsymbol{S}_i \cdot \boldsymbol{S}_j \tag{1.8}$$

where $\boldsymbol{S}_i$ are the spins of the electrons. When $J > 0$, the material will display ferromagnetic ordering, where atomic moments tend to align parallel with each other, as this will lower the exchange energy U. The parallel moments in the material will add up to result in a (often strong) net spontaneous magnetization.

The strength of the exchange interaction J varies between different elements and can be washed out when the thermal energy of the system is comparable with the exchange energy. The temperature where the atomic moments no longer favor spontaneous alignment is known as the Curie temperature (T_C). Above this temperature, the material will be paramagnetic with the susceptibility following the (mean-field) Curie-Weiss law [17]

$$\chi = \frac{C}{T - T_C} \tag{1.9}$$

The reciprocal susceptibility of many ferromagnets is roughly linear in temperature in the paramagnetic regime. Below T_C, the susceptibility diverges with the material showing a finite **M** for zero **H**.

Ferromagnets are the simplest outcome due to an exchange interaction, but materials can have many other types of exchange interactions and behavior. If $J < 0$, the magnetic moments will tend to anti-align with each other, and if the moments of the atoms are all equal, this will result in an antiferromagnetic material. Antiferromagnetism sets in below the Néel temperature (T_N) where the material has zero net magnetic moment. Above T_N the material will be paramagnetic, with (mean-field) susceptibility given as [17]

$$\chi = \frac{2C}{T + \theta} \tag{1.10}$$

Theoretically $\theta = T_N$ and experimentally θ/T_N is on the order of 1, though it can deviate due to longer range interactions.

If the material has more than one species of atom, the interactions can become even more complex. A material with $J < 0$ and two sublattices of different atoms with unequal moments is known as a ferrimagnet. A schematic of different magnetic ordering is shown

in Figure 1.4 and characteristic qualitative susceptibility curves for the different types of magnetism we have discussed are given in Figure 1.5. All of the exchange interactions covered so far are known as “direct exchange” since they are a consequence of direct interaction between the atoms. If a material has magnetic atoms separated by non-magnetic ones, it is possible for the non-magnetic atoms to mediate an exchange interaction between its neighboring magnetic atoms that will dominate the direct exchange between them. This is called “superexchange.” If the exchange interaction is mediated by conduction electrons in metals it is known as “itinerant exchange” [21].

In the preceding discussion I covered the basic atomic origin of different types of conventional magnetism. I now take a step back and consider the aggregate properties of ferromagnetic materials which will be relevant to the study of ASI.

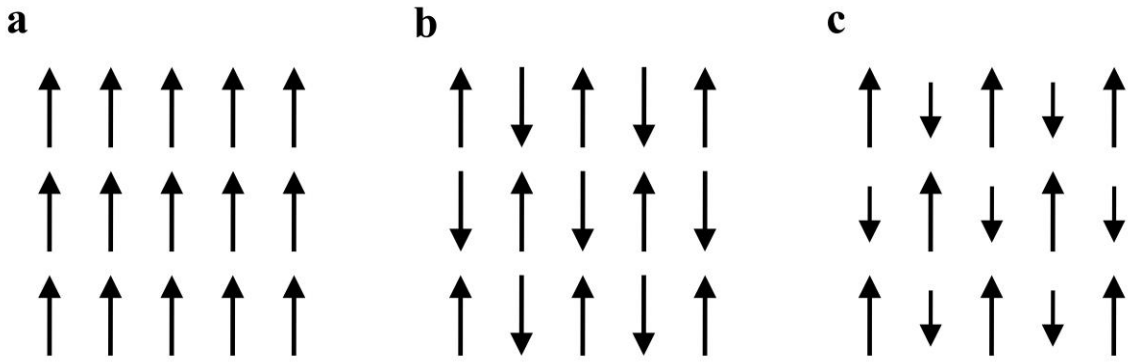

Figure 1.4 **Types of magnetism**. **a**. Ferromagnetic coupling between spins causes them to point in the same direction. **b**. Antiferromagnetic materials have spins of alternating orientation with no net magnetic moment. **c**. Ferrimagnetic materials have magnetic moments that alternate orientation, but one orientation has a higher magnitude so there is still a net magnetic moment.

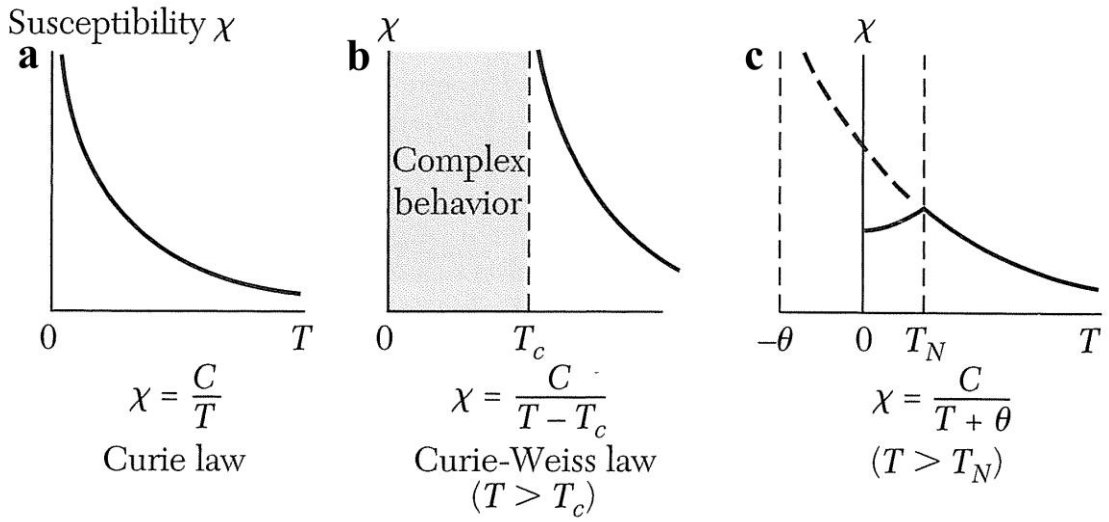

Figure 1.5 **Susceptibility curves**. Magnetic susceptibility vs. temperature for (**a**) Paramagnets, (**b**) Ferromagnets, and (**c**) Antiferromagnets. Also shown are the general equations for susceptibility vs. temperature for different types of magnetism. Figure adapted from Ref. [17].

While it is true that ferromagnetically coupled atoms will align with each other locally, in samples with ~10^{23} atoms, different overall properties can emerge. In considering bulk effects within ferromagnetic materials and between separate ferromagnets, the dipolar interaction will be important. Dipolar coupling exists between any two magnetic moments $\boldsymbol{m}_1$ and $\boldsymbol{m}_2$ and is given by [22]

$$U = \frac{-\mu_0}{4\pi|\boldsymbol{r}|^3}\left(3(\boldsymbol{m}_1 \cdot \hat{\boldsymbol{r}})\,(\boldsymbol{m}_2 \cdot \hat{\boldsymbol{r}}) - \boldsymbol{m}_1 \cdot \boldsymbol{m}_2\right) \qquad (1.11)$$

where $\boldsymbol{r} = \boldsymbol{r}_1 - \boldsymbol{r}_2$ is the vector between the two moments. The dipolar interaction between two neighboring atomic moments is generally much weaker than the exchange interaction, but the exchange energy falls off over a very short range, while the dipolar coupling is a much longer range interaction, decreasing as $\frac{1}{r^3}$. The effect from atomic ferromagnetic moments can add up so that the dipolar interaction becomes significant.

While iron is known to be a very strong permanent magnet, not all naturally occurring pieces of iron will display a net magnetization. This is because a magnetic material can divide into separate magnetic "domains" which can have different orientations. If the domains generally all point in the same direction, the sample will have a spontaneous moment, but if the domains are randomly oriented, the material will not have a strong magnetization. The highest possible magnetization for a sample occurs when all its domains point in the same direction and the maximum value of the magnetization in this configuration is known as the saturation magnetization. The higher the saturation magnetization of a material, the stronger the magnet it can be.

In general, magnetic domains will tend to align with external fields. Consider a sample with a distribution of randomly oriented domains. As a magnetic field is initially applied to the sample, domains preferentially aligned with the field will begin to grow at the expense of unfavorably oriented domains [17]. As the magnitude of the external field increases, all the domains will eventually rotate to align with the field until the sample reaches its saturation magnetization. If the field is removed, the sample will return to a value known as the remnant magnetization (M_r). The applied field value that causes the magnetization of the sample to reach zero after it has been saturated is defined as the coercive field (H_c) and this value will change sign depending on the previous saturation state of the magnet. This phenomenon is known as magnetic hysteresis and a typical hysteresis curve showing the situation I just described is given in Figure 1.6. The shape of the hysteresis curve, or equivalently the values of M_r and H_c, give insight into a magnet's utility for different applications.

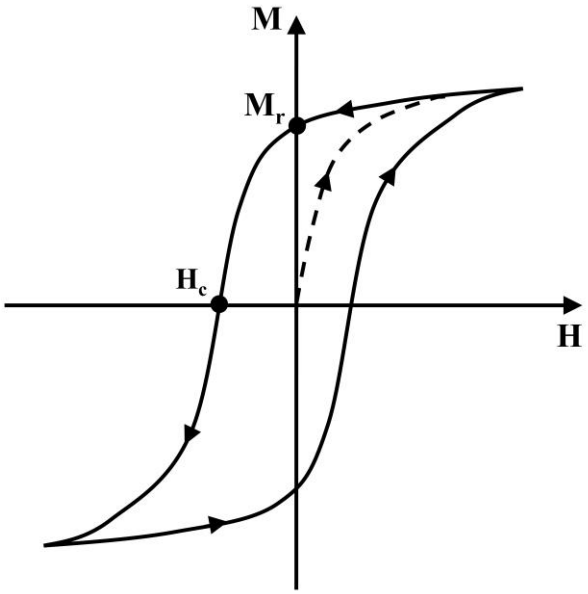

Figure 1.6 **Magnetic hysteresis curve**. Magnetization *M* vs. applied field *H*. An external field is applied to a magnet with randomly oriented domains and no net magnetization. Initially, domains aligned with the field grow in size at the expense of unfavorable domains. As the field is increased, domains will rotate to align with the field until the magnet reaches its saturation magnetization. When the field is removed, the magnetization reduces to the magnet's remnant magnetization, M_r. The field is then applied in the opposite direction. The magnitude required to bring the net magnetization back to zero is known as the coercive field, H_c. The magnitude of the field continues to increase and saturates the magnet in the opposite direction. The value of the magnetization for a given value of applied field depends on the history of the magnet, an effect known as hysteresis.

Some crystals have a preferential axis for magnetization to point along. The preferential axis is called the easy axis and an additional magnetic anisotropy energy arises due to this preference. The crystalline anisotropy originates from spin-orbit coupling creating asymmetric charge distributions around atoms [17]. A preferred magnetic axis will impact the domain configurations through the relative orientation of adjacent crystallographic grains and also through the added anisotropy energy competing with other interactions. It can greatly complicate a magnet's behavior and in general, materials with very low crystalline magnetic anisotropy are better for ASI studies. Shape anisotropy of

the magnet is a separate consideration from crystalline anisotropy and will be discussed in detail below.

Different magnetic domains will be separated by a domain wall where the magnetization switches from one orientation to another. Domain walls often occur at grain boundaries in polycrystalline materials as these are natural locations for a transition, but they don't have to. Neighboring grains can have the same magnetization or two domains can be present in the same, single-crystal grain. In the boundary between two separate domains, the magnetization does not jump discontinuously. There is a characteristic width of the domain wall over which the magnetization will slowly rotate, as this reduces the exchange interaction between the non-parallel spins in the wall. A Bloch wall is most common for bulk magnets, in which the magnetization rotates out-of-plane, as shown schematically in Figure 1.7. In magnetic thin films, if the material is thin enough it may be favorable for the magnetization to rotate in-plane due to external magnetic field. This is defined as a Néel wall and is shown in Figure 1.7(b). In thin-film samples of ASI made from permalloy ($Ni_{81}Fe_{19}$), other types of domain walls, such as a vortex wall, have been observed (see Figure 1.8) [23]. Some examples of magnetic domain patterns and domain walls I observed are shown in Figures 1.9 and 1.10.

Domain walls in thin films have a characteristic width that can be tens to hundreds of nanometers depending on the material and the thickness of the film [24]. When a sample of magnetic material is small enough, it will not be able to stabilize a domain wall and will possess only a single magnetic domain. The exact size required for single domains depends on the material, but it is generally on the order of a few hundred nanometers or less [17]. If the island is elongated, shape anisotropy will cause the magnetization to point along

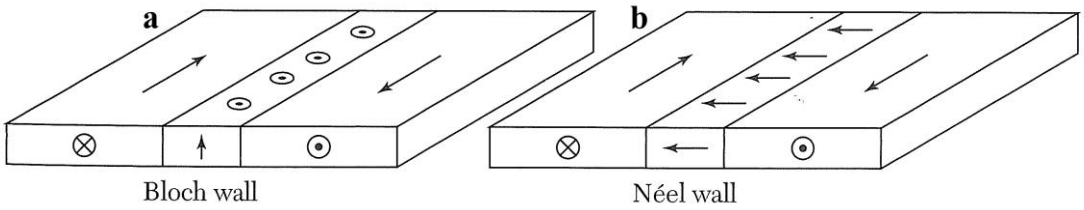

Figure 1.7 **Bloch vs. Néel domain walls**. Relative magnetization orientations for (**a**) Bloch and (**b**) Néel domain walls in a thin film. The magnetization in a Bloch wall rotates out of the plane of the thin film. In a Néel wall, the magnetization rotates in plane. There are different energies associated with the two different types of walls. Figure adapted from Ref. [17].

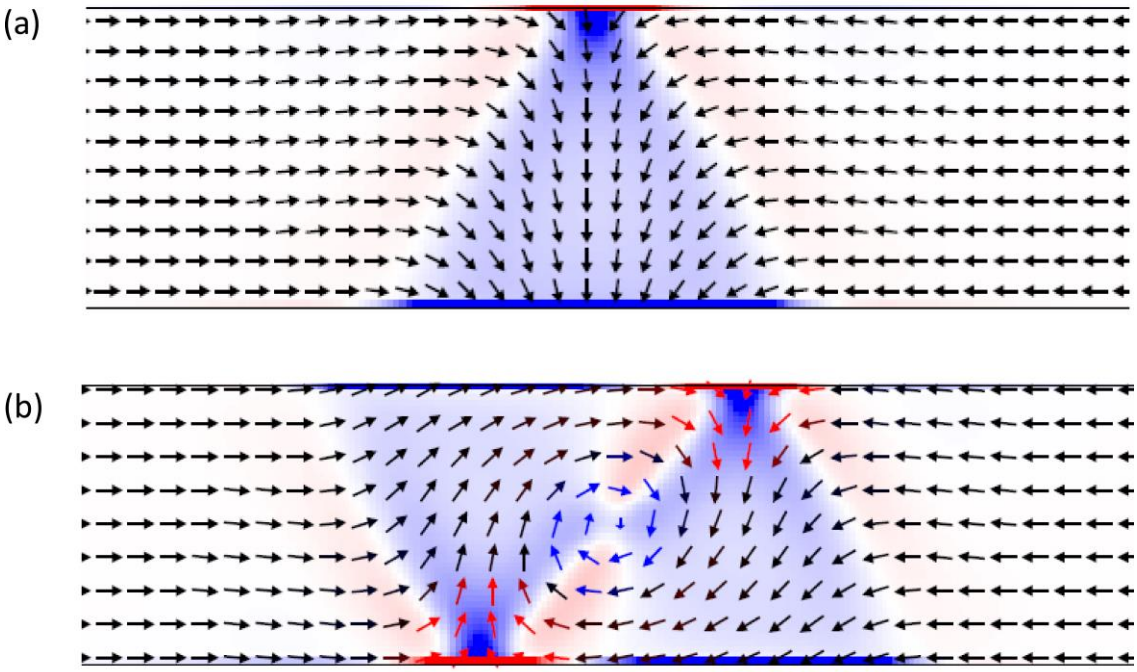

Figure 1.8 **Transverse and vortex domain walls in thin film permalloy**. Micromagnetic simulations of (**a**) transverse and (**b**) vortex domain walls in thin films of permalloy ($Ni_{81}Fe_{19}$). The Transverse domain wall is similar to a Néel wall. The two types of walls are found depending on the thickness of the permalloy. Ladak and coworkers [23] found samples with 18 nm thick permalloy supported transverse walls, while thicker samples made from 36 nm films showed vortex walls. The blue and red color corresponds to positive and negative divergence of the magnetization, which is equivalent to magnetic charge density. Figure adapted from Ref. [23].

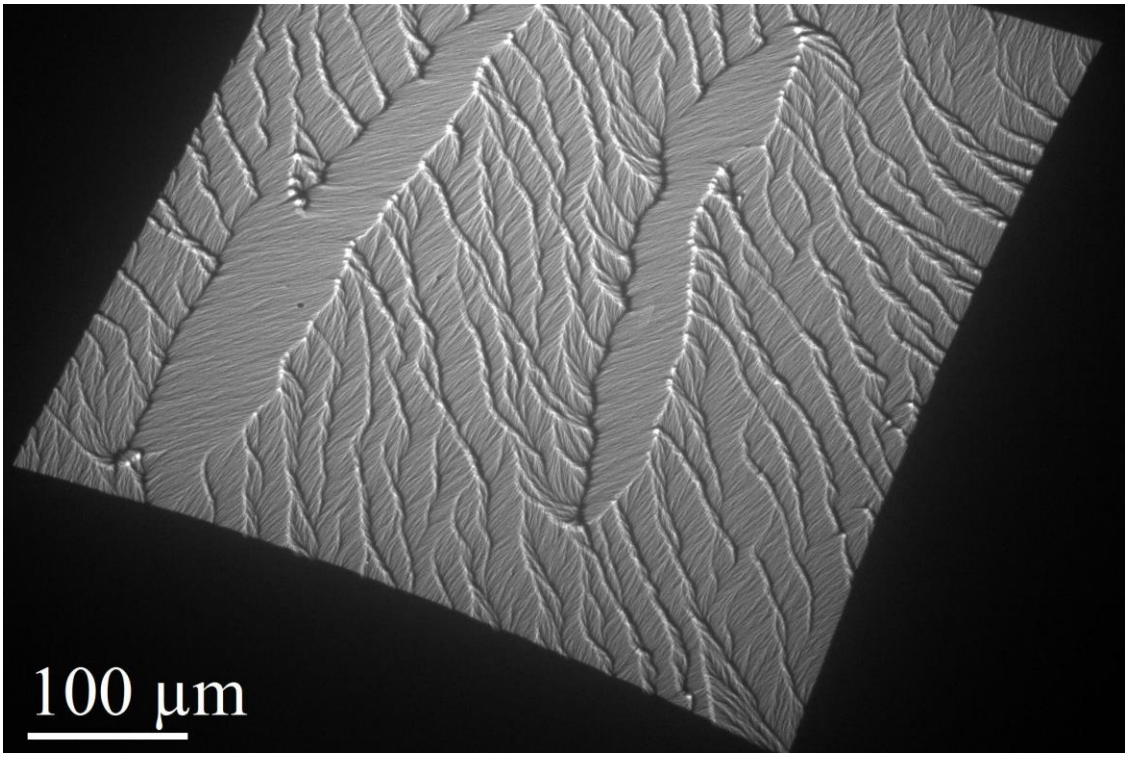

Figure 1.9 **Magnetic domains in thin film $Ni_{35}Fe_{65}$**. Lorentz contrast Transmission Electron Microscopy (TEM) image of a magnetic thin film showing multiple magnetic domains. The thick bright and dark lines are domain walls in the film. The other ripple contrast is representative of small local variations in the magnetic order of the atoms in the material. See Section 2.3.2 for a discussion of Lorentz TEM.

the long axis of the island. This is due to the magnet trying to minimize its external magnetic field energy, proportional to B^2 [22]. Magnetic field lines connect the north and south poles of the single domain magnet by traveling through the free space around it with low density. If the magnetization points along the short axis of an elongated bar magnet, there will be more external flux energy than if the magnetization points along the long axis. Therefore, simple energy minimization can explain the tendency to point in these directions [25]. This is shown schematically in Figure 1.11. A similar consideration explains some domain configurations observed in larger, multi-domain magnetic samples. Neighboring domains will tend to arrange themselves in a configuration that minimizes

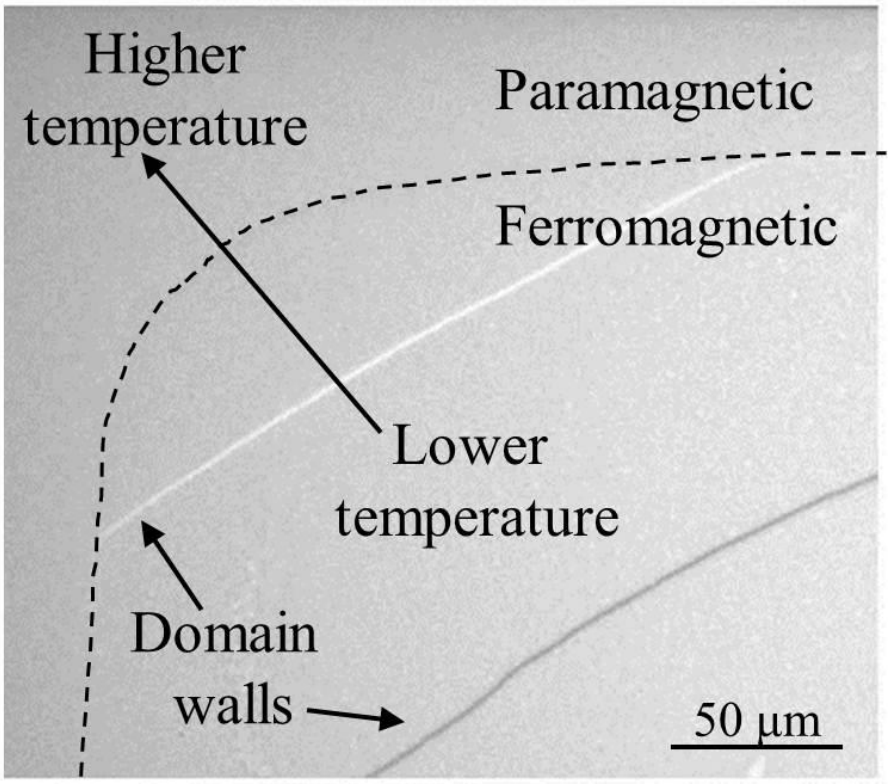

Figure 1.10 **Domain walls and sharp paramagnetic transition in thin film $Fe_{35}Pt_{65}$.** Lorentz TEM image of a thin film heated *in-situ*. A temperature gradient across the film results in both paramagnetic and ferromagnetic regions in the image. The ferromagnetic region has two domain walls visible. See Figure 2.31 for an explanation of the temperature gradient in this image.

the magnetic energy of flux outside of the sample, which often results in circular, or flux closure patterns [17,24]. An example of this is shown in Figure 1.12.

Free spins or magnetic moments can point in any direction, though many times when spins are in a crystal they are confined by lattice geometry to point, *e.g.* in one of only two directions, either "up" or "down." Such two-state spins are called Ising spins and are a good approximation to the way spins behave in many real materials [13,14,26,27]. Ising spins essentially behave as a two level system (they only have two readily accessible quantized energy levels) corresponding to the two spin directions. Other higher energy states may be present, but are typically so much higher in energy that the they can be

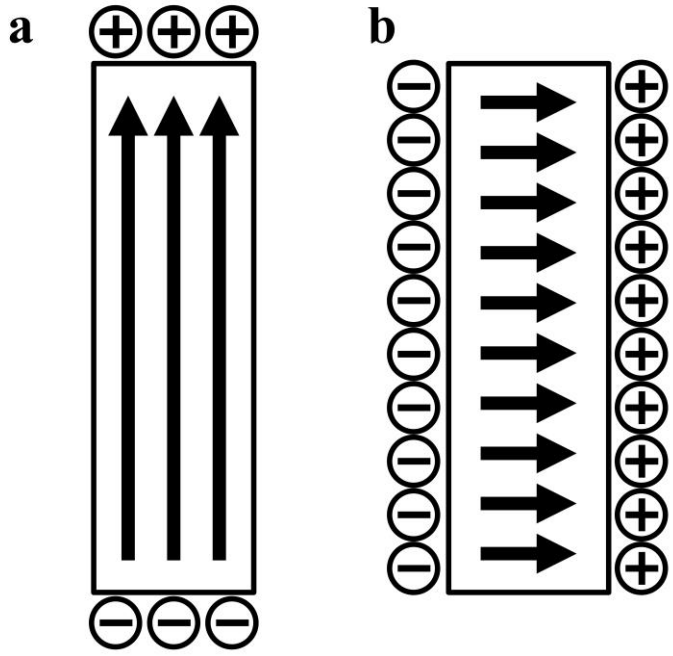

Figure 1.11 **Shape anisotropy flux minimization in single domain magnets**. **a**. When the magnetization points along the long axis of a single-domain, elongated bar magnet, there are fewer magnetic field lines than when the magnetization points along the short axis (**b**). Each effective magnetic charge generates flux lines that connect to an opposite charge on the other side. The greater the area of effective magnetic charges, the more flux lines and the higher the magnetic field energy will be. In general, the bar magnet's magnetization will point along the long axis in order to minimize the field energy.

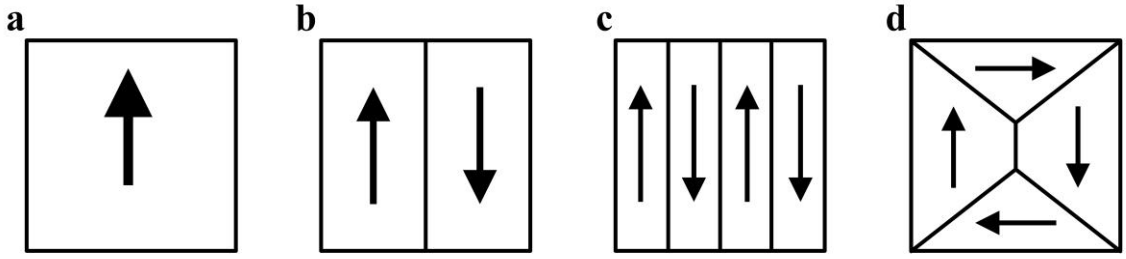

Figure 1.12 **Flux closure domains**. Different domain configurations in multi-domain samples without strong shape anisotropy can help minimize external flux and magnetic field energy. **a**. A single domain has no flux minimization. **b-d**. Domain patterns with decreasing external flux. The flux closure domain configuration in (**d**) has the lowest external flux.

ignored in many situations. Single-domain, elongated, nanoscale bar magnets can act like giant Ising spins that can either point one way or the other along the island's long axis. If Ising-like magnets or magnetic moments are arranged in certain geometries with just the right interactions, it is possible that pairs of spins will not be able to minimize coupling energy with only two states per spin and some moments will be forced into unfavorable arrangements. A situation where all of the pairwise couplings cannot be satisfied results in frustration of the system.

1.2.2 Frustrated magnetism

In magnetic materials, frustration is found when the pairwise interactions between neighboring magnetic moments cannot be simultaneously satisfied. Typically, the frustration arises due to either disorder in the system or its geometry. Prototypical examples are shown in Figure 1.13. Ising spins on a square lattice with ferromagnetic coupling are not frustrated (Figure 1.13(a)). On the other hand, if there is disorder in the material through, for example, impurities or a disordered alloy, this can create varying bond types and lead to frustration (Figure 1.13(b)). Antiferromagnetic exchange between spins on a triangular lattice causes geometric frustration since the geometry of the system does not allow for all bonds to be simultaneously satisfied (Figure 1.13(c)). This is similar to the frustration found in water ice that I discussed in Section 1.1 and is the source of the frustration in the magnetic analog to water ice that I will discuss in Section 1.3, pyrochlore spin ice. Before I explore the rich physics of spin ice, the material that inspired the ASI systems that are the main focus of this dissertation, I describe a few other frustrated spin systems with interesting properties.

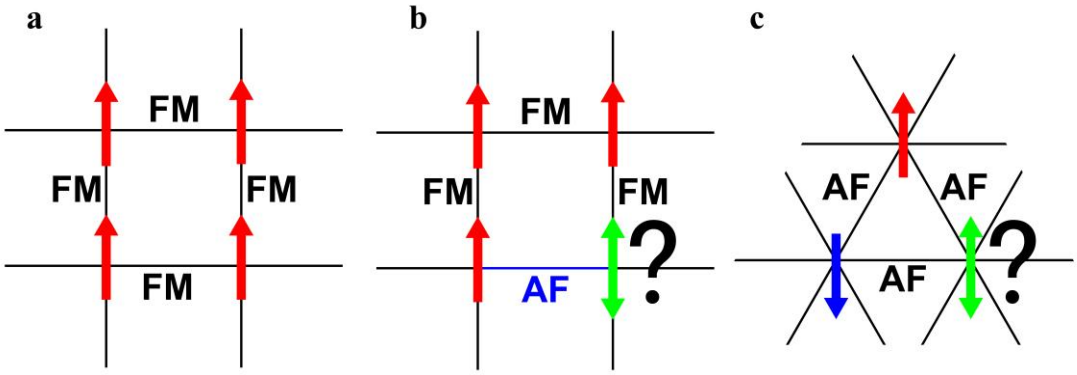

Figure 1.13 **Types of frustration. a**. No frustration: Ferromagnetic interactions between Ising spins on a square lattice are not frustrated. **b**. Disorder frustration: Disordered bonds between Ising spins on a square lattice cannot be satisfied with any spin configuration, leading to frustration. **c**. Geometric frustration: Antiferromagnetically coupled Ising spins on a triangular lattice are frustrated due to the geometry of the system.

Spin glasses are a class of frustrated system where disordered magnetic atoms are randomly arranged among a lattice of non-magnetic atoms [28,29]. The magnetic moments couple through the RKKY (Ruderman-Kittel-Kasuya-Yosida) interaction [30–32] which oscillates (changes sign and magnitude) depending on the distance between the atoms. If the atoms are random distances apart, they have alternating, competing interactions that result in a frustrated system. Well-studied spin glass systems include the dilute magnetic alloy CuMn [33], with Mn concentrations on the order of 1 at. %, and the insulating system $Eu_xSr_{1-x}S$ [34]. Due to random arrangements of atoms and random coupling between the atoms, these systems are frustrated and have difficulty finding an ordered state. There is a fundamental question about spin glasses, as well as other frustrated systems, on whether a unique ground state actually exists. It is predicted that there probably is a unique lowest energy state in any given spin glass, but it is very difficult for the system to find this state. Instead, there are a vast number of low lying energy states that, while not the absolute

lowest energy of the system, are a local minimum in energy in the configuration space. The system will tend to find one of these minimums at low temperature and will not have enough thermal energy to move to a different local or global minimum.

Some materials will display spin fluctuations down to extremely low temperatures [35–38]. $Tb_2Ti_2O_7$ is one such frustrated system. Frustration preventing any long-range order is a fairly typical property, but it is interesting that there are still spin fluctuations down to as low as 70 mK [35], indicating that the energy barrier between degenerate states is extremely small, on the order of 6 μeV. The lack of correlations between spins at any distance is similar to molecules in a liquid and thus this material and others like it are known as "spin liquids." Small perturbations in frustrated systems can actually induce ordered states. In water ice, doping with KOH will cause the lattice to find a unique ground state [39]. Similarly, in $Tb_2Ti_2O_7$, small magnetic fields or applied pressure will result in spin ordering, or "crystallization" of the liquid [36].

Many more complex materials with unique phases and orderings are possible and often times multiple phases and phase transitions will be present in a single material. $Gd_3Ga_5O_{12}$ can display a spin glass phase, spin liquid phase, and antiferromagnetic long-range-ordered phase depending on the temperature and applied field [40]. Neutron scattering experiments on $Zn_2Cr_2O_4$ reveal a spin liquid phase characterized by fluctuations of not individual spins, but clusters of 6 antiferromagnetically coupled spins [41]. Fe_2O_3 nanoparticles display a spin glass phase only in a 0.6 nm thick layer at their surface [42]. These examples only scratch the surface of the vast field of frustrated magnetic materials and there are even more purely theoretical frustrated systems that yield a plethora of new

phases and physics. Overall, the consequences of frustration are diverse, but I now turn my attention back to the materials that motivate the main focus of this dissertation.

1.3 Spin ice

Spin ices are a class of geometrically frustrated magnetic materials with properties that are highly analogous to water ice. The canonical and most frequently studied materials are $Dy_2Ti_2O_7$ and $Ho_2Ti_2O_7$, although other systems such as $Ho_2Sn_2O_7$ also display spin ice physics [13,43–46]. In these materials, the rare earth atoms form a pyrochlore lattice of corner sharing tetrahedra, as shown in Figure 1.14(c). The magnetic moments of the rare earth ions are Ising-like and are confined to point either directly in or directly out of each tetrahedron along the <111> crystal direction. It is believed that the coupling between the rare earth atoms is via a slightly antiferromagnetic superexchange, but ferromagnetic dipolar interactions dominate due to the extremely high atomic magnetic moment of the rare earth atoms [47]. However, the ferromagnetic interactions cannot all be simultaneously satisfied due to the tetrahedral geometry and spin confinement and thus the system is frustrated. The lowest energy configuration for the 4 spins on each tetrahedron has 2 pointing in and 2 pointing out, completely analogous with the ice rules and ordering of hydrogen atoms in water ice. Similarly, there are 6/16 different ways to arrange the 4 spins that satisfy the ice rules resulting in a macroscopically degenerate ground state and residual entropy at zero temperature. The equivalence between spin ice and water ice is shown in Figure 1.14(a) and (b). Ramirez and coworkers were the first to measure the heat capacity of $Dy_2Ti_2O_7$ and calculate the residual entropy, confirming the analogy with water ice [13].

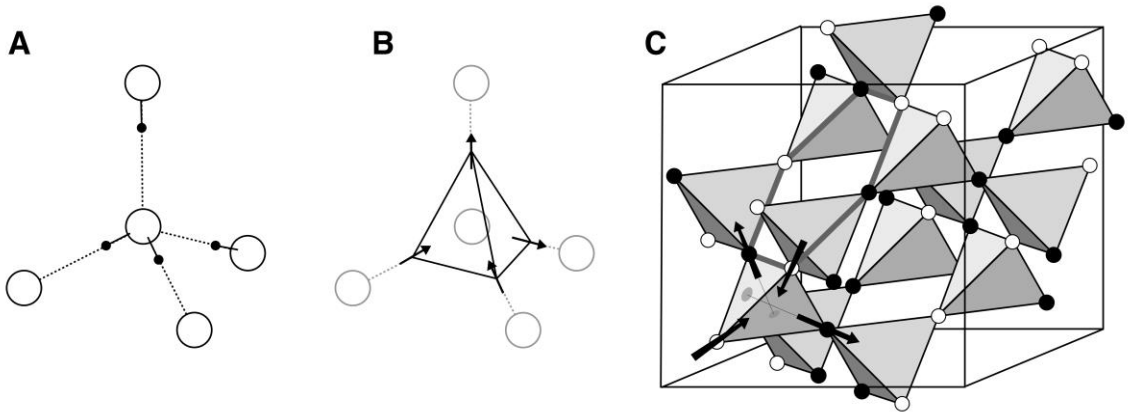

Figure 1.14 **The pyrochlore lattice and the equivalence to water ice**. **a-b**. Two-in, two-out ice rule ordering of hydrogen atoms in water ice maps onto two-in, two-out ordering of rare earth magnetic moments in spin ice. **c**. The pyrochlore lattice of corner sharing tetrahedra showing the ice rule state in spin ice. Magnetic moments are indicated for one tetrahedron and are seen to point directly into or out of the tetrahedron. Black and white dots indicate the spin directions for downward facing tetrahedra and are defined in the tetrahedron with arrows indicated. Figure adapted from Ref. [43].

They successfully match the entropy to a model

$$S_0 = k_B \cdot ln(\Omega_0) = k_B \cdot ln\left(\left(\frac{3}{2}\right)^{\frac{N}{2}}\right) = \frac{N}{2} k_B \cdot ln\left(\frac{3}{2}\right) \qquad (1.12)$$

where the extra factor of ½ comes from mapping N oxygen atoms onto N tetrahedra that contain $2N$ spins in the pyrochlores. Their measurement results are shown in Figure 1.15.

Neutron scattering is a common technique used to gain insight into magnetic correlations within materials [48–50]. The technique is based on the fact that neutrons possess spin, but no charge, so they will diffract in a scattering experiment due to magnetic interactions. Neutron scattering has been used to confirm models for the interactions in spin ice. Figure 1.16 shows experimental neutron scattering data on $Ho_2Ti_2O_7$ compared to simulations based on a nearest neighbor exchange model and a model with long-range dipolar interactions [43,46]. It can be seen that the experimental data matches much better with the long-range dipolar model, confirming the presence of dipolar coupling in spin ice.

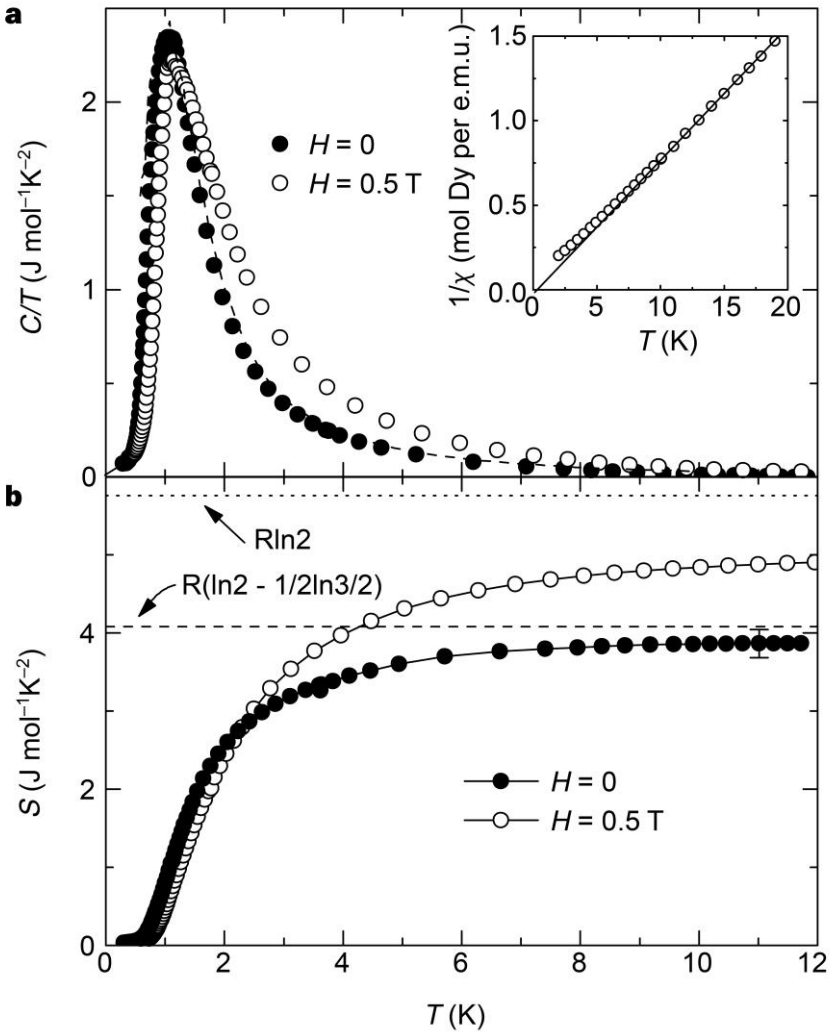

Figure 1.15 **Specific heat and entropy in $Dy_2Ti_2O_7$. a**. Specific heat C divided by temperature T as a function of temperature measurements of $Dy_2Ti_2O_7$ in zero field and H = 0.5 T. **b**. Entropy S as a function of temperature for $Dy_2Ti_2O_7$ in zero field and H = 0.5 T. The entropy is found by integrating the specific heat. The Pauling value of R[ln(2) – 1/2ln(3/2)] is indicated showing the analogy to water ice. Inset: Magnetic susceptibility as a function of temperature for $Dy_2Ti_2O_7$ in an applied field of 0.2 T. Figure adapted from Ref. [13].

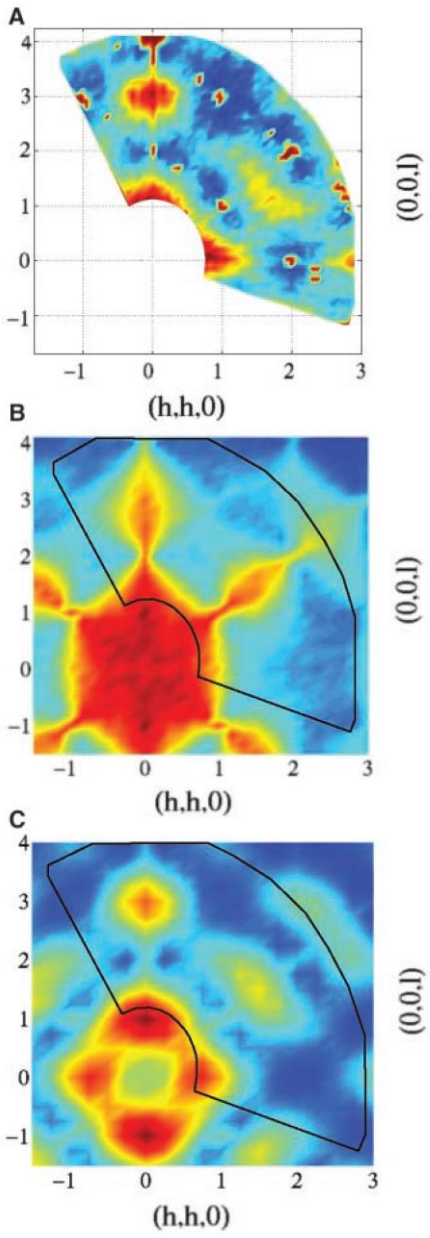

Figure 1.16 **Neutron scattering on $Ho_2Ti_2O_7$**. **a**. Experimental neutron scattering data on $Ho_2Ti_2O_7$. Dark blue is low intensity, dark red is high intensity. **b**. Calculated neutron scattering pattern from a spin ice model with only nearest-neighbor exchange interactions at $T = 0.15J$. Outlined region corresponds to the experimental data in (**a**). **c**. Calculated neutron scattering pattern for a spin ice model with long-range dipolar interactions at $T =$ 0.6 K. This model does a better job of capturing all the features in the experimental data. Figure adapted from Refs. [43] and [46].

Despite measurements confirming a residual entropy in spin ice and implying a disordered and degenerate low-temperature state, a unique ground state has nevertheless been predicted [51,52], due to the long-range dipolar interactions in the system, but it has proven difficult to observe experimentally [13,53–56]. The ground state was predicted by Melko, den Hertog, and Gingras (MDG) and is shown schematically in Figure 1.17 [51]. The state can be described by an order parameter and associated ordering vector $\boldsymbol{q} = (2\pi/a)<001>$. MDG predict a jump in heat capacity and associated transition to the ground state with no residual entropy at 0.18 K. Recently, a single crystal sample of $Dy_2Ti_2O_7$ thermally equilibrated for extremely long times (> 100 hours) showed a slight upturn in heat capacity and a reduction in entropy, below the Pauling value [57]. Additionally, a thin film sample of $Dy_2Ti_2O_7$ with a high level of strain was also reported to find an ordered state and recovery its entropy, though this is different than the MDG ground state [58]. The MDG ground state has yet to be confirmed in spin ice and I will give some insight on a potential explanation for its elusiveness and propose some promising ways to reach it in Chapter 4.

One property of the spin ice system that has attracted great interest is its excitations above the spin ice state. The spin ice state has perfect 2-in, 2-out ordering of magnetic moments in each tetrahedron, consistent with the ice rules. If we consider each spin as a dumbbell of magnetic charge, then this state has two positive and two negative charges in each tetrahedron and these add up to no net charge (Figure 1.18(c)). Excitations above this manifold are spin flips that create pairs of neighboring tetrahedra with 3-in, 1-out and 3-out, 1-in arrangements of spins. In the dumbbell picture, these elementary excitations create pairs of magnetic monopoles, regions of excess magnetic charge inside the

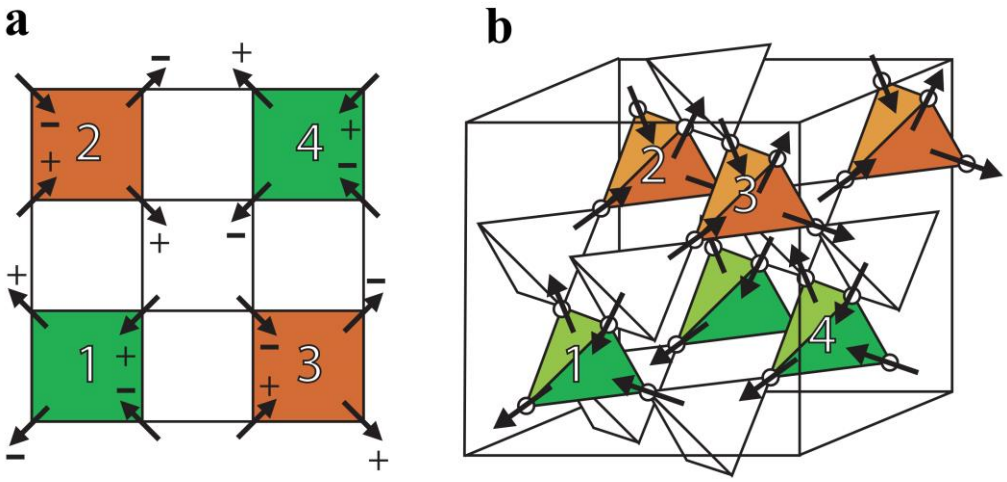

Figure 1.17 **Spin ice ground state. a**. Predicted MDG spin ice ground state projected down the z axis. +/- indicates the component of each spin in the z direction. **b**. Three dimensional schematic showing one unit cell of the ground state. Figure adapted from Refs. [43] and [51].

neighboring tetrahedra (Figure 1.18(d)). Once the monopoles are created, it only costs a Coulomb-like energy for them to separate and move throughout the lattice, leaving behind a Dirac string (Figure 1.18(e)) [59].

Monopoles in spin ice were first proposed by Castelnovo *et al*. in 2008 as emergent phenomena and the elementary excitations in spin ice [59]. Shortly thereafter, monopole signatures were observed experimentally [55,56,60,61] and the monopole model was successfully used to explain previously poorly understood magnetic relaxation measurements on $Dy_2Ti_2O_7$ [62]. Neutron scattering experiments further confirmed the monopole description [61] and the presence of Dirac strings in these systems [55]. Muon spin relaxation measurements reported a magnitude of the monopole charge to be 5 μ_BÅ^{-1} [63], confirming a prediction in Ref. [59], although the validity of these results has been challenged [64,65]. Still, the monopole picture has proven to be an excellent theoretical model for describing the physics of spin ice and as a method to interpret and understand

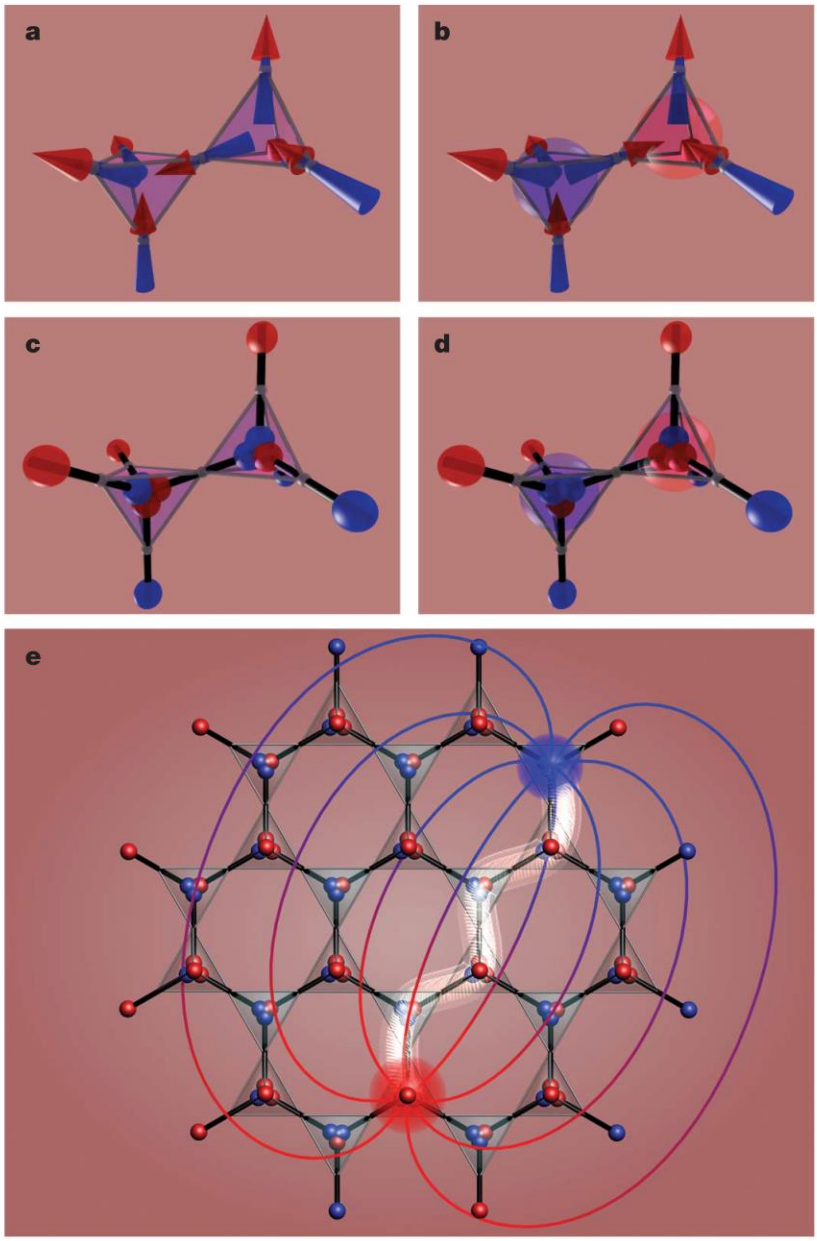

Figure 1.18 **The dumbbell picture and magnetic monopoles in spin ice**. The images in (**c**) and (**d**) correspond to replacing the spins in (**a**) and (**b**) with dumbbells where the ends are opposite magnetic charges. Frames (**a**) and (**c**) show the spin ice state which has two positive and two negative charges in each tetrahedron and thus no net magnetic charge. Frames (**b**) and (**d**) show an excitation of the system, where one spin is flipped, creating neighboring tetrahedra each with excess positive or negative magnetic charge, or monopoles. Once the monopoles are created, it costs no more energy other than the standard Coulomb interaction for them to separate and move throughout the lattice. The monopoles will leave behind a Dirac string along the path they take, indicated in frame (**e**). Figure adapted from Ref. [59].

experimental data on these systems [66,67]. Some groups have taken the analogy to electricity even further, coining and proposing terms and ideas such as “magnetolyte” for a medium to carry magnetic charge, “magnetronics” for circuits utilizing magnetic charge, and “magnetricity” for monopole motion in a magnetic field as an analog to how electrons, electric monopoles, move in an electric field [63,68].

It is important to note that monopoles in spin ice are monopoles in **M**, not **B**. Dirac predicted the existence of monopoles in the magnetic **B** field in the 1930’s to help explain the quantization of electric charge [69]. Despite a great number of experiments and searches for them [70–73], no isolated Dirac monopoles have been ever been confirmed. One study at Stanford University did see a single potential monopole event, although this was never duplicated, independently verified, or confirmed [70]. Dirac monopoles were recently observed in a synthetic magnetic field, although this does not indicate that they exist naturally [74]. Though the monopoles in spin ice behave in the same manner as Dirac’s proposed monopoles in **B**, their magnitude is different than a Dirac monopole and they do not explain the quantization of the charge of an electron. They can be thought of more as quasi-particle excitations that describe the physics of spin ice at low temperatures, similar to other particle-like descriptions of physical phenomena such as photons or phonons [21,22]. Dirac strings in spin ice are an observable phenomenon [55], whereas Dirac strings of **B** in the vacuum have no reported experimental signature. Overall, the pyrochlore spin ices are an intriguing class of materials that display rich physics and novel phenomena originating from the extensive frustration in the system.

1.4 Artificial spin ice

1.4.1 Introduction

In the study of frustrated systems such as ice and spin ice, it would be ideal to have the ability to image exact atomic and spin configurations to better understand the effects of frustration and the dynamics in the system. Unfortunately, this is not possible with current technology. An alternative approach to studying these systems is through lithographically patterned, two-dimensional arrays of Ising-like, single-domain nanoscale bar magnets, collectively known as artificial spin ice (ASI) [75,76]. The individual bar magnets behave like pseudo-atoms or giant "spins" that interact in a classical manner. ASI arrays are highly tunable systems built by design to mimic the frustrated interactions in real materials, but are also amenable to a variety of direct characterization techniques such as Magnetic Force Microscopy (MFM) [75,77–79], Photoemission Electron Microscopy (PEEM) [80–85], Magneto-Optical Kerr Effect (MOKE) [86,87], x-ray resonant magnetic scattering [88,89], or Lorentz Transmission Electron Microscopy (LTEM) [90–95]. LTEM is the technique I use and will discuss in depth in Section 2.3.2. ASI presents a playground for studying frustrated systems, where the geometry of the system as well as the strength and character of the magnetic interactions are readily chosen by design. Precise information at the single spin level gives new and unique insight into frustrated systems and is a huge advantage of ASI over other experimental systems and techniques.

Through this work, I aim to answer some long-standing and fundamental questions about frustrated systems. Namely, how do these systems find long-range ordered states or is this even possible at all in the presence of frustration? In frustrated systems, an apparent

degeneracy of ground states prevents long-range order, often when detailed analysis of perturbations predict that an ordered state nevertheless should occur [51]. Despite decades of intense interest, frustrated systems still pose fundamental problems, with many unanswered questions, due in part to the tendency of these systems to inefficiently explore their configuration spaces and to lose ergodicity [96]. Monte Carlo simulations can address some of these issues through the introduction of more complicated basic excitations [97,98], but questions about the specific conditions that will allow or restrict a system from reaching a predicted ground state often go unanswered in a general sense. An emerging approach seeks to address these issues by fabricating physical simulators of real materials [99–101], and ASI comprise a highly productive example.

Materials-by-design approaches to condensed matter physics have recently produced noteworthy results, exemplified by photonic crystals and materials with a negative index of refraction [102,103]. Metamaterials are a growing trend, where "bottom up" methodology allows one to realize physical structures not found in nature in order to investigate and understand genuinely new properties and phenomena [104,105]. These materials systems provide unique, tunable control over parameters previously frozen in or set by a single choice of material or stoichiometry. They have allowed for the realization of old science fiction ideas like invisibility through optical cloaking by precise control over electromagnetic fields [106–108]. An artificially patterned lattice recently demonstrated the first skyrmions at room temperature without the application of a magnetic field [109]. Certain types of arrays of Superconducting QUantum Interference Devices (SQUIDs) are metamaterials which exhibit unique microwave transmission properties [110,111]. Clearly,

materials-by-design systems are a fruitful route to new ideas and new physics and as we will see, artificial spin ices are an excellent example.

1.4.2 Square and kagome ice

The first ASI geometry studied was square ice [75], chosen for its analogy to spin ice with 4 spins that meet at each vertex. An overview of the geometry is shown in Figure 1.19. The lowest energy states follow the two-in, two-out ice rules, although in square ASI the 6 ground-state configurations split into two energy levels due to the inequivalent interaction strengths between the four magnets at each vertex. The variation in the distance between adjacent and opposite spins at each vertex lifts the six-fold degeneracy of the ice rule state and results in 2 configurations having slightly lower energy than the other 4. The 16 different possible spin arrangements are traditionally grouped by their energy and given a type I-IV nomenclature. The lowest energy, type I configurations have two spins pointing in and two pointing out with zero net magnetic moment. Type II vertices also obey the ice rules, but have a net magnetic moment and slightly higher energy than type I. Type III vertices are three-in, one-out or one-in, three-out arrangements, and type IV are four-in or four-out. Schematics of these vertex types are shown in Figure 1.19(d).

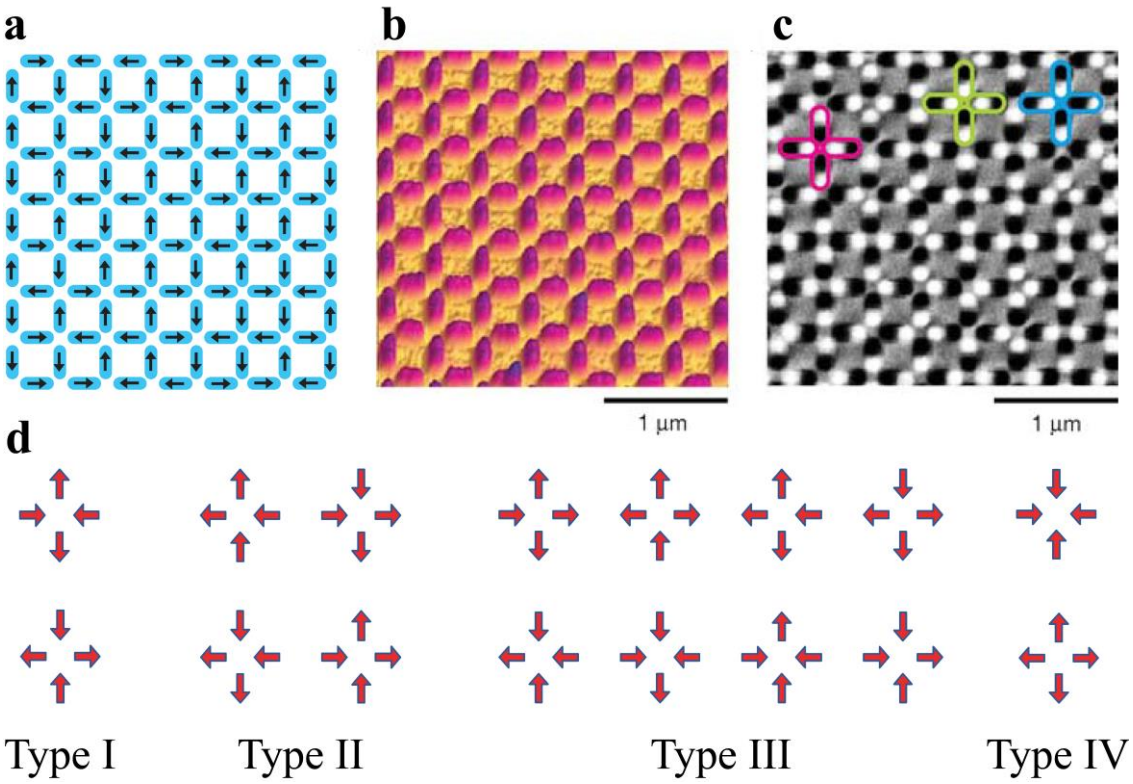

Figure 1.19 **Artificial square ice. a**. Schematic of artificial square ice geometry. The arrows represent the directions of the spins in part (**c**). **b**. Atomic Force Microscopy image of a square ASI sample. **c**. Magnetic Force Microscopy image of a square ASI sample. The moment directions can be exactly determined and are shown in part (**a**). The outlined regions in pink, blue, and green correspond to vertices of type I, II, and III respectively. **d**. The 16 possible moment configurations for a single vertex, grouped into the 4 types based on their relative energy level. Figure adapted from Ref. [75].

The ground state of artificial square ice is a perfect tiling of type I vertices [112]. This state is two-fold degenerate, owing to the two spin configurations that make up the type I energy level. These two ground states are related to each other by a global reversal of spins. Much of the early work in ASI used rotation demagnetization protocols in an attempt to minimize array energy and access ground-state configurations [75,113–116]. Rotation demagnetization is a process where the sample is rotated in a magnetic field that is alternating in polarity at some frequency and also slowly decreasing in amplitude [113]. A schematic of typical protocols is shown in Figure 1.20. The seminal first ASI paper by

Wang *et al*. used rotation demagnetization on square ice samples with varying lattice spacings [75]. This resulted in random populations of vertex types for the larger lattice spacings and ice rule ordering of spins below a random population for the shorter lattice constant samples. Closer magnets increase the magnetostatic coupling between nearest neighbors and favor more ordered states. Indeed, the ordering is seen to emerge as the lattice spacing is decreased as in Figure 1.21. Still, even for the shortest lattice spacing, the samples remain far from the true ground state [75]. Rotation demagnetization was generally not very successful at accessing large regions of ground-state ordering or long-range ordered states [75,114–116], though a recent report using a linear decrease in applied field showed that by tuning the rotation rate of the sample as well as the rate at which the field is ramped down, moderately sized domains of ground-state order can be generated [117].

Disorder is also known to play an important role in energy minimization and how well square geometry samples can find ordered states [118–121]. Specific effects of disorder and its origins in both the square and kagome lattices will be discussed later in this section. In general, other experimental techniques beyond rotation demagnetization are needed to realize long-range ordered states and these methods will be explored in depth in the next section and later chapters.

The spin configurations of artificial square ice are typically frozen in place at room temperature, but the statistical populations of vertex types still yield interesting physical results. Different demagnetization protocols and lattice spacings (coupling strengths) lead to different relative populations of vertex types in a given sample. Using these statistical ensembles, a framework for effective thermodynamics and temperature has been developed by Nisoli and coworkers [115,116]. The probability that a system will be in a given state

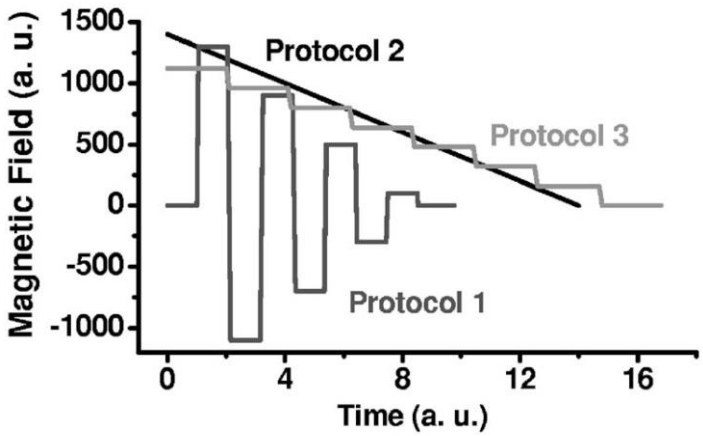

Figure 1.20 **Rotation demagnetization protocols**. Three possible rotation demagnetization protocols used in ASI studies. The amplitude of the magnetic field changes at a given rate and is reduced to zero. Figure adapted from Ref. [113].

is described by the Boltzmann distribution and this probability is a function of the temperature, energy, and number of possible states of the system [122]. Since the exact microstate of an ASI sample can be precisely characterized and the total number of states fully enumerated, the Boltzmann distribution can be worked backwards to yield an effective temperature for a given a sample. More ordered states, closer to the ground state, have lower effective temperatures and other thermodynamic quantities, such as the configurational entropy, can be extracted as well [78,115,123].

The artificial kagome lattice has also been extensively studied in both theory and experiment [78,93,95,124,125]. The kagome lattice is relevant and experimentally useful since it is derived from the pyrochlore lattice found in the rare earth spin ices. The plane of atoms perpendicular to the <111> crystal axis in the pyrochlore lattice of corner sharing tetrahedra is a kagome lattice of corner sharing triangles (see Figure 1.22(a)). Applying a magnetic field along the <111> crystal direction will isolate kagome planes since those spins will not couple to the external field. In artificial kagome ice, only three nanomagnets

meet at each vertex and so the ice rules become “two-in, one-out” and “one-in, two-out” arrangements of spins (see Figure 1.22(b)). The local lowest energy configuration is six-fold degenerate, just like the pyrochlore spin ice state, and there are two higher energy vertices corresponding to three-in or three-out spin arrangements.

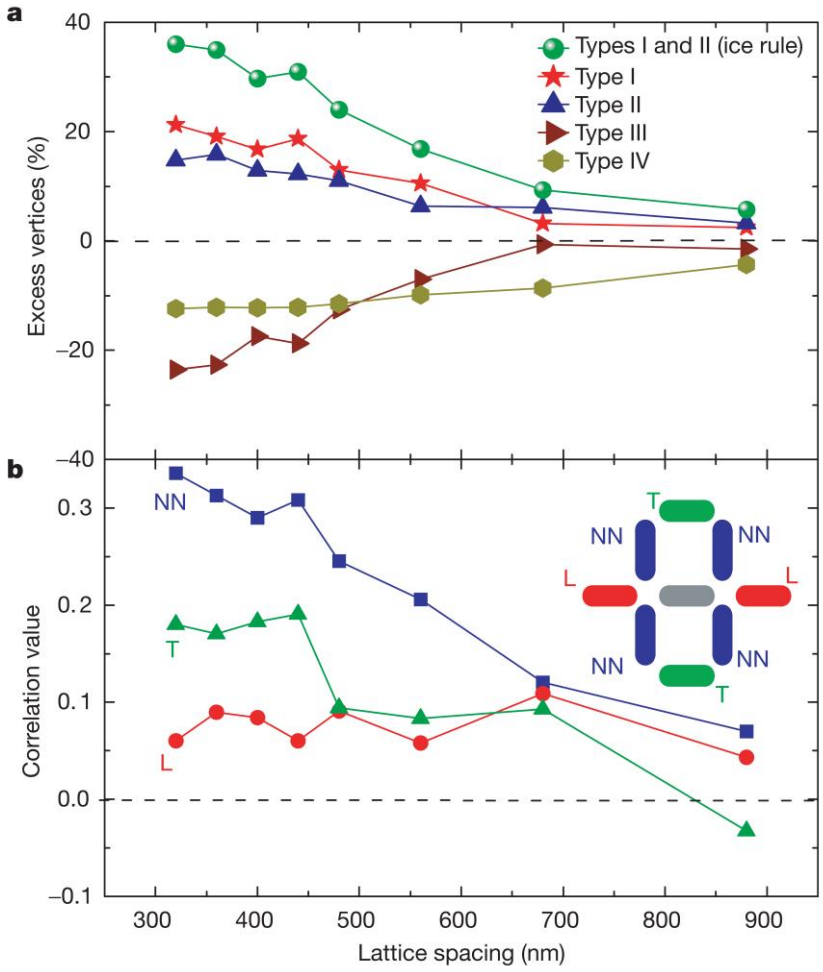

Figure 1.21 **Statistics of moment configurations after a demagnetizing protocol**. **a**. Deviation of vertex populations from a random configuration as a function of lattice spacing between neighboring moments. For larger spacings, the population is nearly random. As the spacing is decreased and the magnetostatic coupling between the spins is increased, the populations become no longer random and transition to a nearly ice-rule state. **b**. Average moment correlation values for different neighbors as a function of lattice spacing. If neighbors are in a favorable (unfavorable) configuration, the pair is assigned a 1 (-1). The average over all pairs in the lattice is given as the correlation value. Figure adapted from Ref. [75].

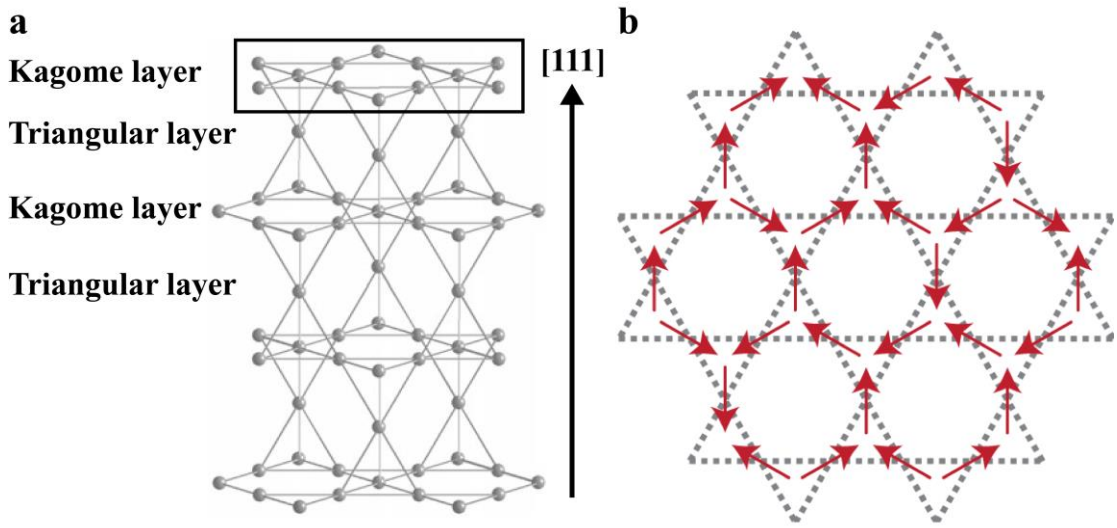

Figure 1.22 **Artificial kagome ice**. **a.** A plane of the pyrochlore lattice of corner sharing tetrahedra perpendicular to the [111] crystal axis is a kagome lattice of corner sharing triangles. A kagome plane is outlined at the top. These kagome layers alternate with triangular layers in the pyrochlore lattice. **b.** A kagome lattice with spins added to each lattice point showing ice rule ordering. The spins can either point directly in or directly out of each triangle and follow the one-in, two-out or two-in, one-out ice rule ordering. Figure adapted from Ref. [126].

Rotation demagnetization experiments on kagome ice yield nearly equal populations of the six ice-rule vertex types [93]. While the ice-rule state in this lattice is easily achieved due to the extensive degeneracy of the vertices, artificial kagome ice is predicted to have a rich phase diagram with multiple ordered states and phase transitions due to long-range dipolar interactions in the system [98,125,127]. The calculated phase diagram is shown in Figure 1.23. For the highest effective temperatures, the system is disordered, paramagnetic, and does not obey the ice rules. As the temperature is lowered, the first phase transition and plateau in entropy corresponds to the ice-rule state, known as the kagome ice I state. This is still an extensively disordered state. The next phase transition corresponds to a charge-ordered state, which is known as the kagome ice II state. As in Figure 1.18, ASI bar magnets can be thought of as dumbbells of magnetic charge with the

magnitude $|Q| = \frac{\mu}{l} = \frac{MV}{l} = MA$, where μ is the magnetic moment of the element, given by the magnetization *M* times the volume *V*. *l* is the length of the magnet or dumbbell and *A* is then the cross sectional area of the nanomagnet [22]. For ice-rule obeying states, there will be an excess +*Q* or –*Q* charge at each vertex (see Figure 1.24(a)). Configurations that violate the ice rule (Figure 1.24(b)) have +3*Q* or –3*Q* charges and are highly energetically unfavorable. The charge-ordered state has alternating +*Q* and –*Q* charges on neighboring vertices and a considerably lower entropy than the ice-rule state (see Figure 1.24(c)). There is an even further ordered state within the charge-ordered manifold which displays long-range spin order and is known as the loop state [98,125,127]. This state has zero entropy as seen in Figure 1.23 and is shown schematically in Figure 1.24(d).

Both the square and kagome ASI lattices can be used to study magnetic monopole excitations analogous to the ones found in the pyrochlore spin ices. Unlike in the real crystalline materials, the monopoles and their interactions in ASI can be imaged directly for a more precise understanding of their properties and behavior. The ground state of the square lattice is a perfect tiling of type I vertices. If one spin is excited and flipped within the ground-state order, it will create a pair of neighboring, oppositely charges monopole excitations as two type III vertices. The three-in, one-out and one-in, three out pair of vertices have +2*Q* and –2*Q* charges, respectively. These monopoles can separate and move throughout the lattice, leaving behind a Dirac string of excited type II vertices connecting them [76]. In the kagome lattice, there are ± *Q* charges at each vertex already and the monopole excitations are the ± 3*Q* ice-rule violating vertices. Monopoles have been proposed as the mechanism for magnetic reversals in ASI structures and have been directly

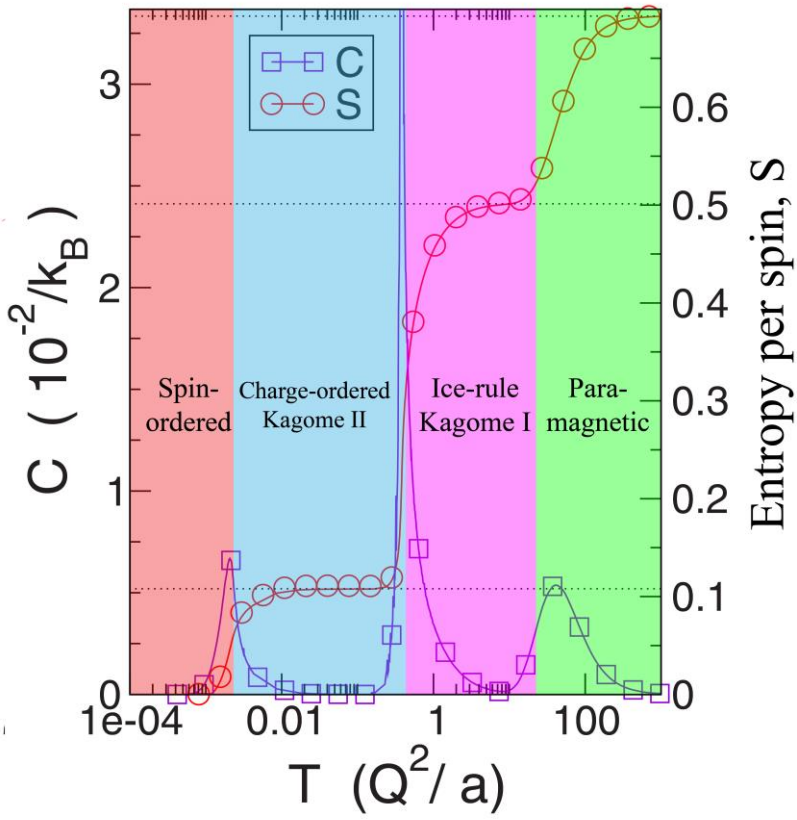

Figure 1.23 **Phase diagram of artificial kagome ice**. Calculated phase diagram showing specific heat *C* (blue) and entropy per spin *S* (red) of kagome ASI as a function of effective temperature *T*. At high temperatures, the system is paramagnetic. As the system is cooled it first orders into an ice-rule obeying state known as the kagome I state. As it is cooled further, there is a predicted charge-ordered state known as kagome ice II. This state is shown in Figure 1.24(**c**). Finally, there is a transition to a spin-ordered state within the charge-ordered manifold. This state is shown in Figure 1.24(**d**). Figure adapted from Ref. [127]**.**

observed in both the square (see Figure 1.25) and kagome (see Figure 1.26) geometries [75,80,94,128].

Disorder is known to play an important role in experimental investigations of ASI where it can arise from a variety of sources [94,95,118–121,130,131]. ASI samples are generally patterned using electron beam lithography (see Section 2.2.1). While this technique, produces very precise and controlled arrangements and is an essential

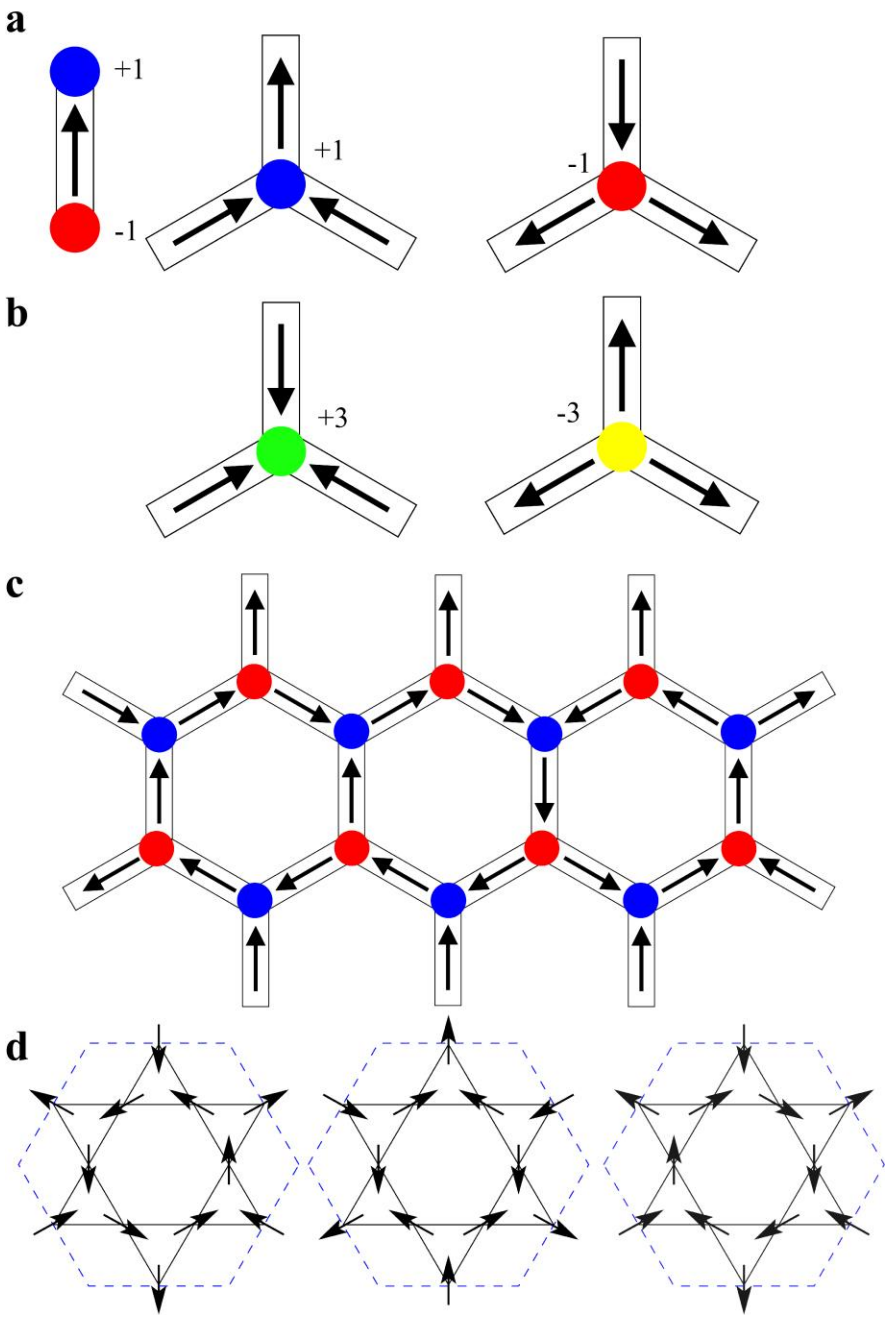

Figure 1.24 **Dumbbells of magnetic charge and ordered states in artificial kagome ice. a.** ASI element modeled as a dumbbell of magnetic charge and schematics of ice-rule obeying moment configurations. Ice-rule obeying states in the kagome lattice have an excess $+Q$ or $-Q$ charge at each vertex. **b.** Schematic of excited, ice-rule violating states in the kagome geometry. Ice-rule violating states have an excess $+3Q$ or $-3Q$ at each vertex. **c**. Schematic of charge-ordering in kagome ASI. The predicted charged ordered kagome II state has alternating $+Q$ and $-Q$ charges on neighboring vertices. **d**. Schematic of the spin-ordered, zero-entropy ground state in kagome ASI. There are six possible ground-state tilings. Three are shown here and the other three are given by a global reversal of spins. Part (**d**) adapted from Ref. [98].

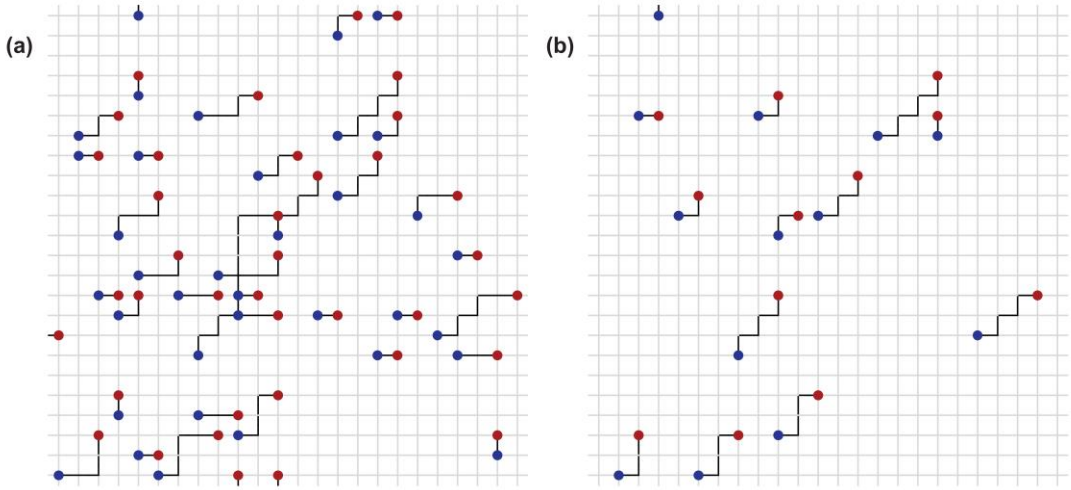

Figure 1.25 | **Magnetic monopole motion in square ASI**. Monopoles in the square lattice, type III vertex excitations connected by a Dirac string, are seen to move in the application of an applied field from (**a**) to (**b**). Figure adapted from Ref. [129].

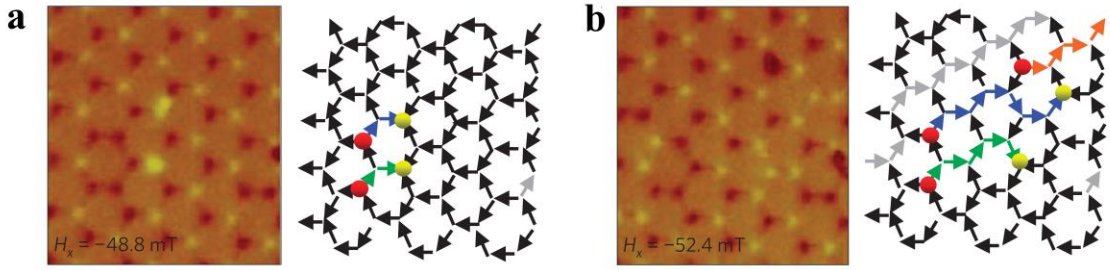

Figure 1.26 **Direct observation of magnetic monopoles in kagome ASI**. MFM images and corresponding spin maps of a kagome ASI sample in an applied magnetic field. Two pairs of ice-rule violating monopoles are observed to move apart from each other in the applied field from (**a**) to (**b**), leaving behind Dirac strings highlighted in green and blue. Another monopole also appears in (**b**) with its partner out of the field of view. Figure adapted from Refs. [80] and [128].

fabrication tool, is still subject to some variation [75,95]. Magnets can have a distribution of widths or differences in edge roughness due to shot noise in the electron beam during lithography. There can be height variations in the material deposited for the magnets and different amounts of surface roughness depending on the deposition technique. Magnets may have slight variations in their location, not being placed in a perfectly ordered array

[121]. These factors relating to the size of the magnets affect the coupling between neighboring ASI spins because they change the magnetic moments of the spins and thus the strength of the dipolar interactions between them. Besides variations in coupling strength, there can also be differences in the coercive field for individual magnets originating from variations in shape, size, or local microstructure in the magnetic material [94,121,132].

While the origins of disorder in ASI are diverse, the effects are surprisingly consistent. Budrikis and coworkers showed in the square lattice that while there are different sources of disorder, their general effect is to broaden the energy levels of the different vertex types [121]. This is shown in Figure 1.27 where, as the amount of the different types of disorder is increased, the energy of each vertex spreads into a range of possible energies where eventually some of the ranges overlap, blurring the energetic distinction between vertex types. It makes ground-state ordering less energetically favorable and makes it easier for systems to access higher energy vertices and vertex configurations. Defects and disorder are also known to cause pinning sites for domain walls. I have observed this in my experiments and Monte Carlo simulations, as discussed in later chapters (see, for example, Figure 4.46).

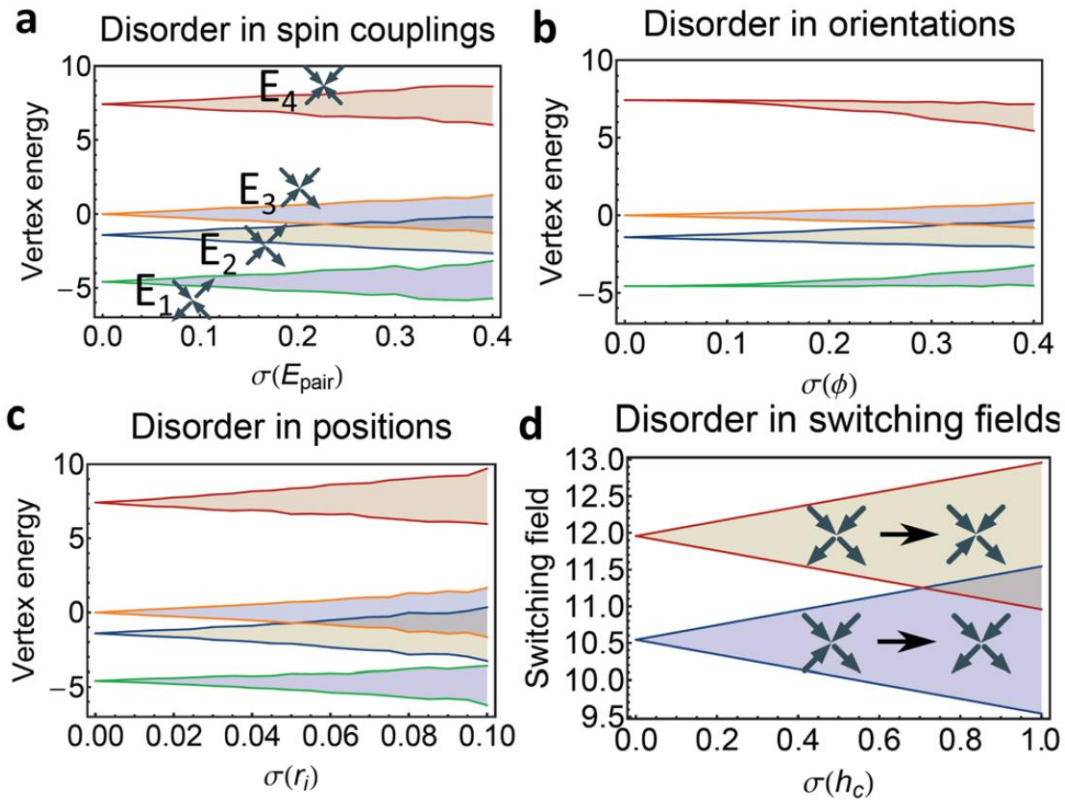

Figure 1.27 **Disorder in square ASI**. Disorder from different sources has similar effects on the energy landscape of the vertex types in square ASI. From top to bottom, the graphs show the relative energy of type IV, III, II, and I vertices due to disorder in (**a**) pairwise energy, (**b**) orientation, and (**c**) position. **d.** Disorder in the field required to switch a type II to a type III vertex and vice-versa. Figure adapted from Ref. [121]**.**

An interesting set of experiments performed in the Cumings group by Stephen Daunheimer on the kagome geometry gives additional insights into disorder in ASI systems. Daunheimer's work involved applying a strong magnetic field to polarize kagome samples along one direction, rotating the sample 100°, 120°, or 180°, and then slowly increasing the magnetic field from zero until the sample was completely polarized in the other direction, capturing the state of the system at each field step. Daunhimer found that the individual magnets began to flip in stronger applied fields and the total magnetization of the crystal slowly reversed. This technique provides a measure of the disorder in the

system by giving the spread in coercivities for the individual magnets and by the functional form of the magnetization curve. An example of these measurements is shown in Figure 1.28. The shape of the reversal curves is different for different rotation angles due to the relative orientation between the applied field and the three magnetic sublattices in the kagome geometry. Depending on the field angle, different magnetic sublattices are more (or less) aligned with the applied field and couple more (or less) strongly to it. This technique can also be used to isolate the effects of the separate individual sublattices on the total magnetization so the sublattices can be studied independently. The experimental data can be fit with error functions so that the average coercivity and characteristic spread in coercivities can be extracted from the fit. The ratio of the spread in coercive fields to the average coercive field defines a figure of merit for the disorder in a sample. Daunheimer showed that fabricating kagome samples out of permalloy and in a connected geometry reduces the disorder compared to prior studies of kagome ASI [80,94,128].

A noteworthy feature of the 180° magnetic reversals of artificial kagome ice is avalanches of reversed magnets. Avalanche phenomena are seen in a wide variety of physical systems including sand piles, rice piles, earthquakes, and Wolfram's cellular automata [133–137]. A close inspection of the magnetization curves in Figures 1.28 and 1.29 will reveal significant discontinuous jumps in the data. These jumps are large chains, or avalanches, of magnets that all reverse together in a single field step. The jumps are random and can take on many shapes, as seen in Figure 1.30. The avalanches are different every time, suggesting that there is stochastic disorder in the system [132]. The sizes of the avalanches follow a power law distribution with a cutoff given by the size of the finite system, or correspondingly, the largest possible avalanches (see Figure 1.31). Power law

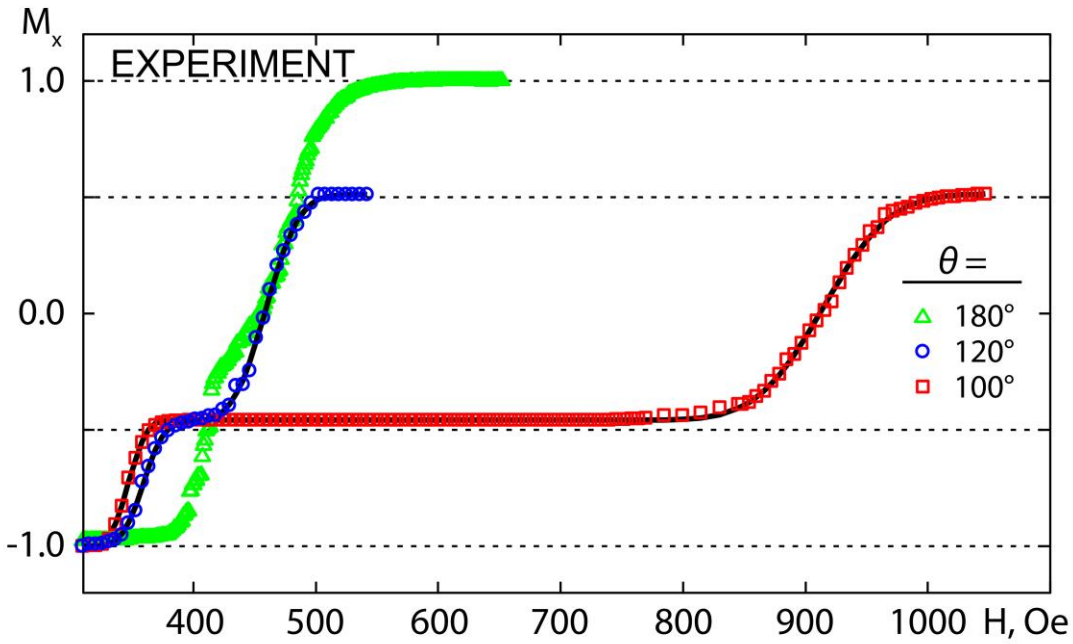

Figure 1.28 **Disorder in kagome ASI**. Magnetization vs. applied field for magnetic reversals of kagome ASI samples in three directions relative to the original polarization direction along one sublattice. The shape of the magnetization curves are a superposition of two error functions which fit characteristic spreads in the coercivities of individual sublattices that reverse as the applied field is increased. The spread in coercivity is a measure of the disorder in the system. The curves differ for the different reversal directions because the various sublattices couple more or less strongly to the applied field depending on the reversal angle. Figure adapted from Ref. [94]**.**

behavior is characteristic in other systems that have avalanches, such as granular media and earthquakes, and power laws that stretch over several orders of magnitude are reflective of the self-organized criticality of the system [138].

In all of the pioneering work on ASI, experiments and samples were exclusively athermal. That is, they showed no changes in their properties or behavior as a function of temperature. Samples were only studied by applying global, external magnetic fields. However, applying arbitrary magnetic fields only allows for certain transitions between different states of the system. In order to access predicted long-range ordered states, such as the ground state in the square lattice or the charge-ordered state in the kagome geometry,

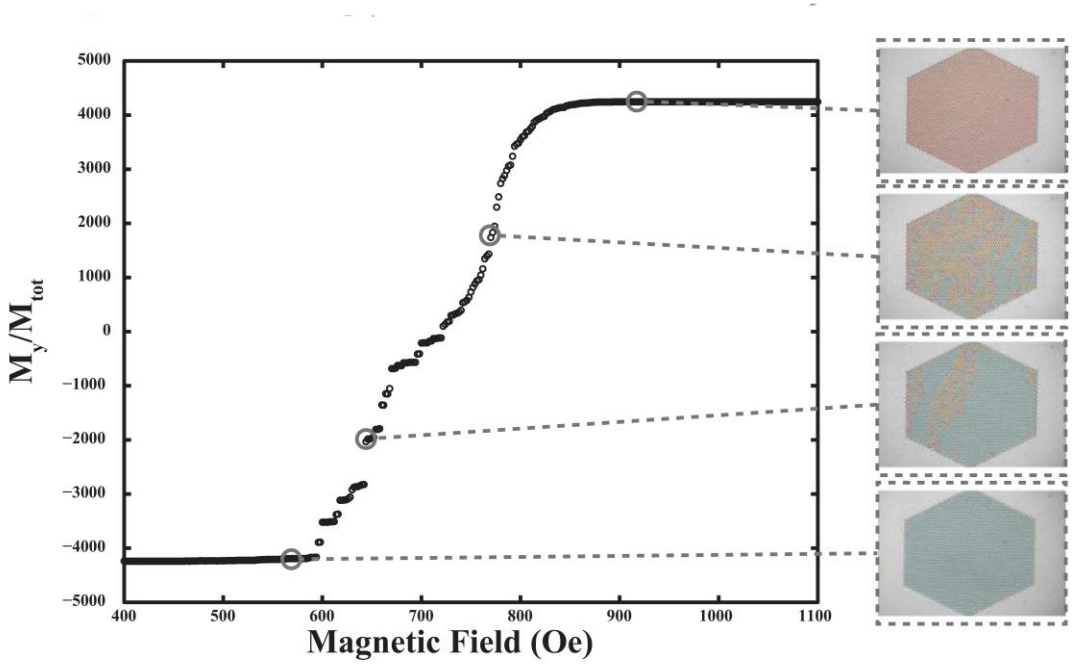

Figure 1.29 **Avalanches in artificial kagome ice**. Measured Magnetization vs. applied field in a 180° reversal of a kagome ASI sample. The large jumps in the magnetization data are due to large avalanches of magnets that all reverse at the same time. On the right are snapshots of the crystal at different points in the reversal process. Figure adapted from Ref. [132]**.**

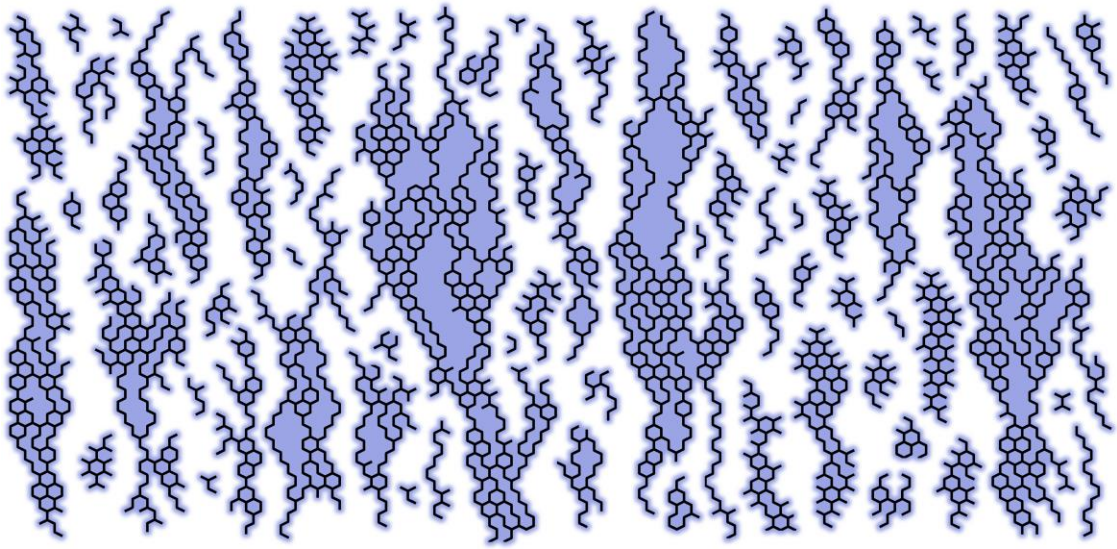

Figure 1.30 **Shapes of avalanches**. Representative example of some of the shapes of avalanches observed by Daunheimer in his kagome ASI magnetic reversal experiments.

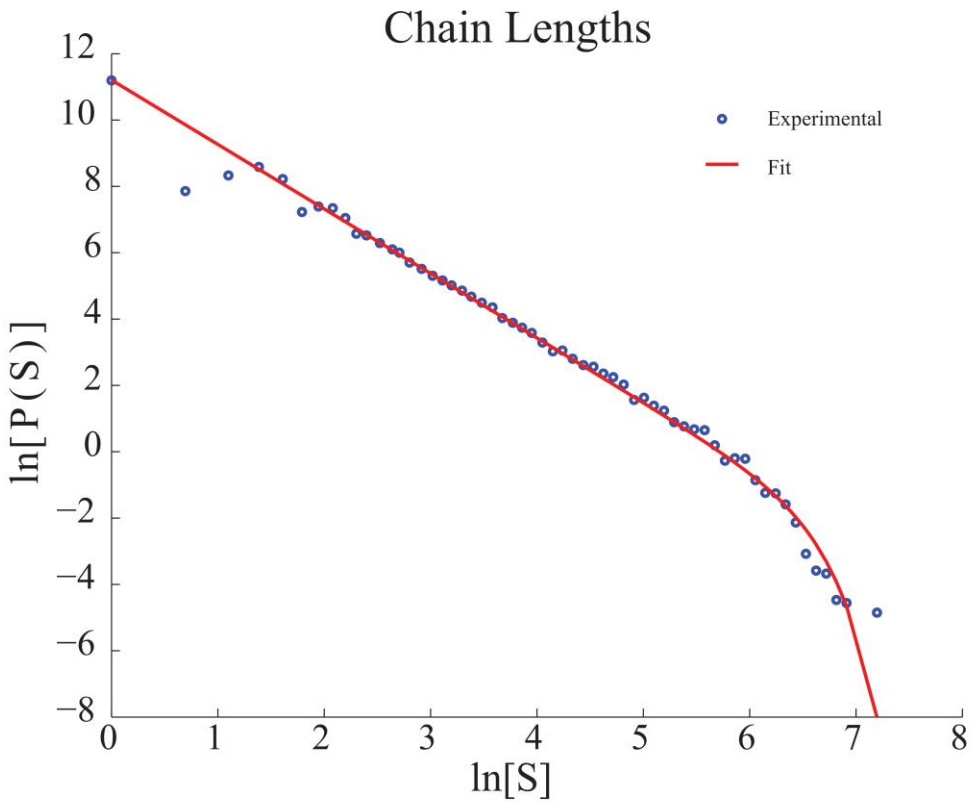

Figure 1.31 **Avalanche distribution in artificial kagome ice**. Power law distribution of avalanche sizes in magnetic reversals of kagome ASI samples. Log of the probability of observed avalanche size vs. avalanche size. Experimental data is fit to a power law with a cutoff related to the size of the finite crystal studied. Figure adapted from Ref. [132].

more transitions between different states of the system need to be opened up. One approach to achieving this is to allow for thermally activated reversals of individual magnets.

1.4.3 Thermal activation

Ice and spin ice have properties that explicitly depend on and vary greatly with temperature. Spins and atoms are constantly interacting and fluctuating, even down to the lowest temperatures. To improve artificial spin ice as a model for real systems and to expand on its functionality for investigating long-range ordered states in frustrated systems,

thermal activation of the system has been introduced [78,79,81,82,84,85,87,95,139–141]. Including temperature dependent effects vastly increases ASI's capability as an experimental tool. It allows for deep investigations of ordered and low energy states, as magnets that are allowed to interact and flip on their own will naturally try to minimize their energy within their local environment [78,79,85,95]. Allowing the system to relax and lower its energy makes it much more analogous to other physical systems that tend to find ground states. Thermal activation also introduces dynamics into the system – real timescales which can be measured and implemented in simulations to give a more realistic picture of what will happen in experimental systems [81,82]. This section will discuss techniques used to introduce thermal effects into ASI and give a survey of the noteworthy results thus far. The field is rapidly growing and this dissertation's main focus in later chapters is on improving and adding to the techniques and knowledge of thermally active ASI.

Before discussing the techniques and results of thermal flipping of magnets in ASI, it will be beneficial to understand the physics behind why magnets spontaneously flip magnetization. An elongated bar magnet has an anisotropy energy density K and volume V, giving it a total anisotropy energy KV. This energy must be overcome in order to flip the magnet, so if the thermal energy in the system is raised close to the anisotropy energy, the magnet will flip at a rate given by an Arrhenius law [142]

$$\tau^{-1} = \nu_0 exp(\frac{-KV}{k_B T}) \tag{1.13}$$

where v_0 is a constant dependent on the material, KV is the anisotropy energy, k_B is Boltzmann's constant, and T is the temperature [143]. We can see that when KV is comparable to k_BT that the rate for the magnet to flip approaches v_0. Microscopically, what

is happening is thermal fluctuations are overcoming the exchange interaction energy that defines ferromagnet order, causing the magnet to lose a defined moment. The atoms may then reorder with a new magnetization direction which can be different from the initial direction. When the magnet is flipping often enough that it is hard to even define a magnetization, this is known as superparamagnetism and was first explained by Néel [144].

In a nanomagnet without magnetocrystalline anisotropy, the anisotropy energy KV becomes the shape anisotropy energy of the magnet [81]. Magnets that interact with each other, such as in ASI arrays, are influenced by their local magnetic environment which also affects the flip rate. In this situation, KV is replaced by $E_0+\Delta E$, where E_0 is the shape anisotropy energy, or the intrinsic energy barrier of the magnet and ΔE is the magnetostatic energy change if the magnet were to flip. The shape anisotropy energy is given by [145]

$$E_0 = \frac{1}{2} V D_Z M^2 \tag{1.14}$$

where V is the volume, M is the magnetization, and D_z is the demagnetizing factor along the long axis of the magnet. If we consider an isolated magnet, ΔE is zero and we can see that if E_0 is small, the flip rate for the magnet will be high. This can be accomplished with either a small volume, as in a very thin magnet, or a low magnetization. The intrinsic energy barrier for a permalloy nanomagnet is on the order of 10^4 K and dipolar coupling between the nanomagnets is also on the order of 10^4 K. This is why thermal fluctuations were not seen in the initial ASI experiments.

A key point in considering the flip rates of nanomagnets is that the Curie temperature (T_C) sets an energy scale for the system. As a magnetic material approaches T_C from below, the saturation magnetization reduces continuously until it reaches zero at T_C. A small magnetization will greatly reduce E_0 and ΔE, allowing for thermal flipping. If

the temperature of the magnet is below $T_{C,}$ it will also still have a finite magnetization and be able to interact with neighboring magnets. Essentially, when the thermal energy of the system is close to, but not above T_C, the magnets will flip spontaneously in order to minimize their local energetic configuration. The temperature of the onset of thermal flipping is known as the blocking temperature (T_b) [144]. It is also well known that thin films of a magnetic material can have a lower T_C than the bulk form [24,78].

A seminal study by Morgan and coworkers found large domains of ground-state ordering in the square lattice could be formed through thermalization occurring during the growth of the ASI magnetic nanostructure [77]. As the magnetic material is deposited and the individual magnets grow in thickness, their volume and T_C increases and the nanostructures go through a regime where they are superparamagnetic and able to flip spontaneously due to their temperature and local magnetic environment. This allows the magnets to lower their local energy configurations. By controlling the deposition rate and the substrate temperature, Morgan *et al*. allowed the system to order during growth. Once the magnets become thick enough, the blocking temperature becomes too high and the magnetization is frozen in place. The magnetic configuration was imaged immediately after fabrication (see Figure 1.32). This technique was successful at revealing the long-sought square ground state, but the approach does not allow for much control over the samples or process. In order to efficiently and completely answer fundamental questions about these materials, more robust and flexible techniques are needed, allowing for more control over the system, the material parameters, the nanostructure geometries, the heating processes, as well as the ability to perform repeated experiments.

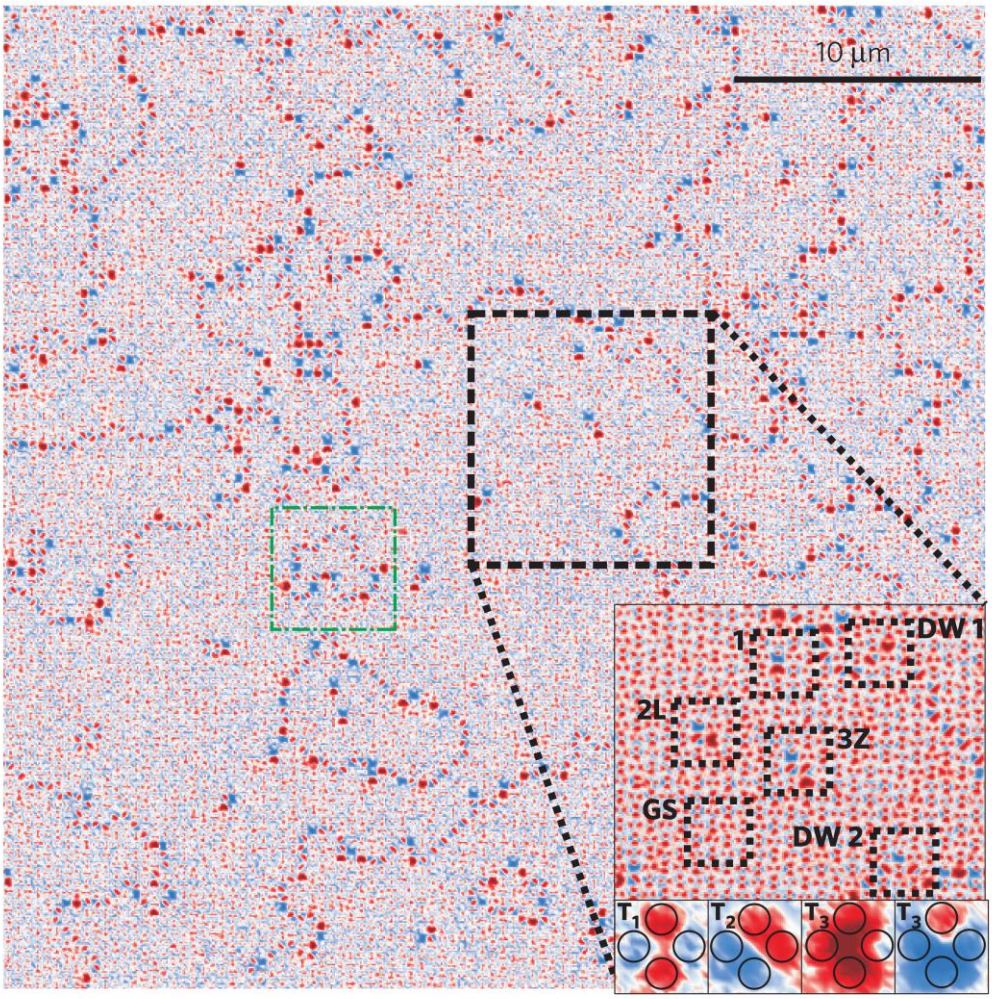

Figure 1.32 **Ground-state domains formed during growth**. MFM image of a square ASI sample that thermalized during growth. Large regions of ground-state ordering of type I vertices separated by higher energy vertex domain walls. Inset shows vertex types and how they appear in the false color MFM contrast. Figure adapted from Ref. [77].

More recently, a new experimental approach has been devised that seeks thermal activation in ASI by heating close to the T_C of the constituent materials [78,79,81,82,84,87,95,139–141]. In this technique, nanoscale magnets begin to exhibit spontaneous reversal, driven simultaneously by thermal fluctuations and by the local magnetic environment. However, the success of this approach thus far has been limited by the choice of material and the nanoislands' geometries. Most studies have used permalloy

($Ni_{80}Fe_{20}$), which has a relatively high T_C, approximately 800 °C in the bulk material [146]. To accommodate such a high ferromagnetic ordering temperature, realizations utilize either very narrow temperature ranges, which are limited at the upper end by degradation of the nanostructures [78,79], or very thin films (≈3 nm) [81,82,85], which simultaneously reduce the onset temperature of thermally activated reversal and the desired magnetostatic interactions of the nanoscale bar magnets. Another technique is to use δ-doped Pd(Fe) monolayer stacks [84,87], which allow for a highly tunable T_C, but the films are still very thin (∼7 nm) and do not allow for strong magnetic coupling between macro-spins. A CoGd alloy has also been used to investigate the kagome lattice [140,141]. It has a relatively low T_C, around 200 °C, though the material is ferrimagnetic and the magnetization density is still significantly lower than that of the commonly-used permalloy. Section 2.1 of this dissertation is dedicated to developing a material which overcomes these challenges and tradeoffs.

Despite the challenges with these techniques, work along these directions has produced some exemplary results using them. Kapaklis and coworkers first demonstrated that the net magnetization of a patterned array of δ-doped Pd(Fe) ASI decreased before the T_C of a continuous film of the same material, suggesting the nanomagnets were spontaneously flipping [87]. Later that year, Arnalds *et al*. showed thermally active building blocks of kagome ASI – one, two, or three hexagon rings – that ordered into ground-state or low energy configurations [84]. For this they used they also used δ-doped Pd(Fe) monolayer stacks. This system was then investigated further, describing the full energetic landscape of the few-hexagon samples and measuring the dynamics of the samples with PEEM, this time by fabricating arrays from ultrathin permalloy [81]. The

material was deposited using a moving shutter method, so elements on one side of the sample where thinner than on the other side. The researchers then looked for regions of the sample that were fluctuating around room temperature and used those crystals for their studies. These magnets turned out to be ~3 nm thick, so they flip readily near room temperature, but don't couple to each other as strongly as other systems. The same group also published results that directly showed the relaxation process from a polarized state to a nearly perfectly ordered ground state in the square geometry [82]. Snapshots of the process are shown in Figure 1.33. It is interesting to note how the samples relaxed towards the ground state. One can see chains of magnets that reverse initially, followed by the rest of the magnets filling in the space between the chains.

Shortly thereafter, a study on both the square and kagome lattice was published by another group who used thicker permalloy nanoislands, ≈25 nm high, heated to ≈550 °C [78]. This work also showed large regions of ground-state order in the square geometry, with elements spaced closer together resulting in larger continuous ground-state domains. More importantly, Zhang and coworkers provided the first experimental demonstration of the charge-ordered kagome ice II state, shown in Figure 1.34. Similar to the square arrays, the size of the kagome charge-ordered domains increases when the spins are spaced closer together. However, the observed charge-ordered domains are fairly small and still quite far from the predicted perfectly charge-ordered state and the heating experiments with thicker permalloy are quite difficult. The substrate for these studies needs special treatment to reduce lateral diffusion of the permalloy. Still, there is only a very small temperature window where the experiments can be performed. Zhang reports magnets begin to flip at

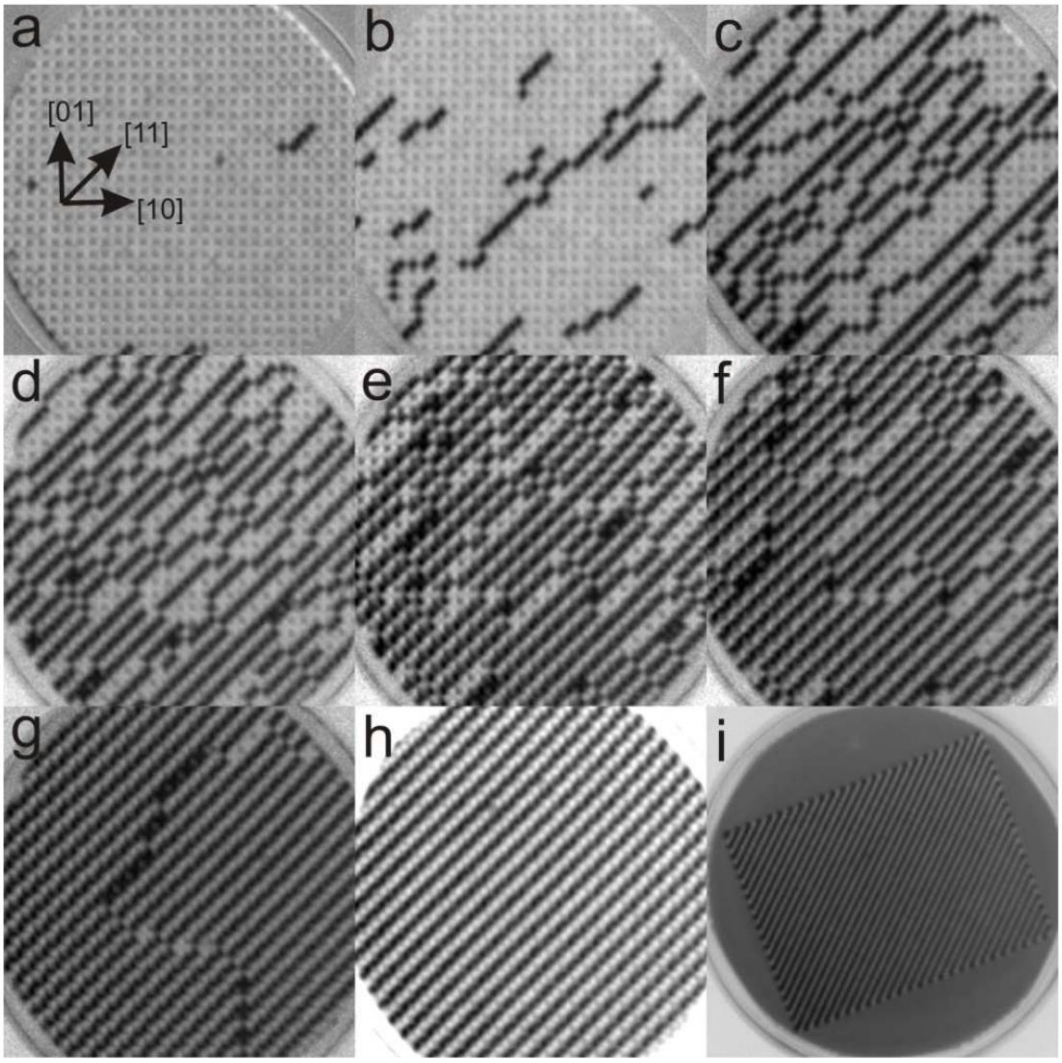

Figure 1.33 **Direct observation of relaxation in square ice**. PEEM with x-ray magnetic circular dichroism (XMCD) images of thermal relaxation in square ASI made from ultrathin permalloy. The sample first relaxes with "strings" and then forms large domains of ground-state order. Field of view in (**a**)-(**h**) is 20 μm, field of view in (**i**) is 50 μm. Figure adapted from Ref. [82].

545 °C and structural degradation of the magnets sets in at 560 °C. This makes repeated experiments difficult and samples cannot last very long. Clearly, there are challenges and tradeoffs with all these methods for thermal activation of ASI. In Chapter 2, I introduce a new material for realizing thermal effects in ASI which will greatly improve the ease of

these types of experiments, allow for more thorough explorations of the parameter space of these systems, and aid in answering fundamental questions about frustration.

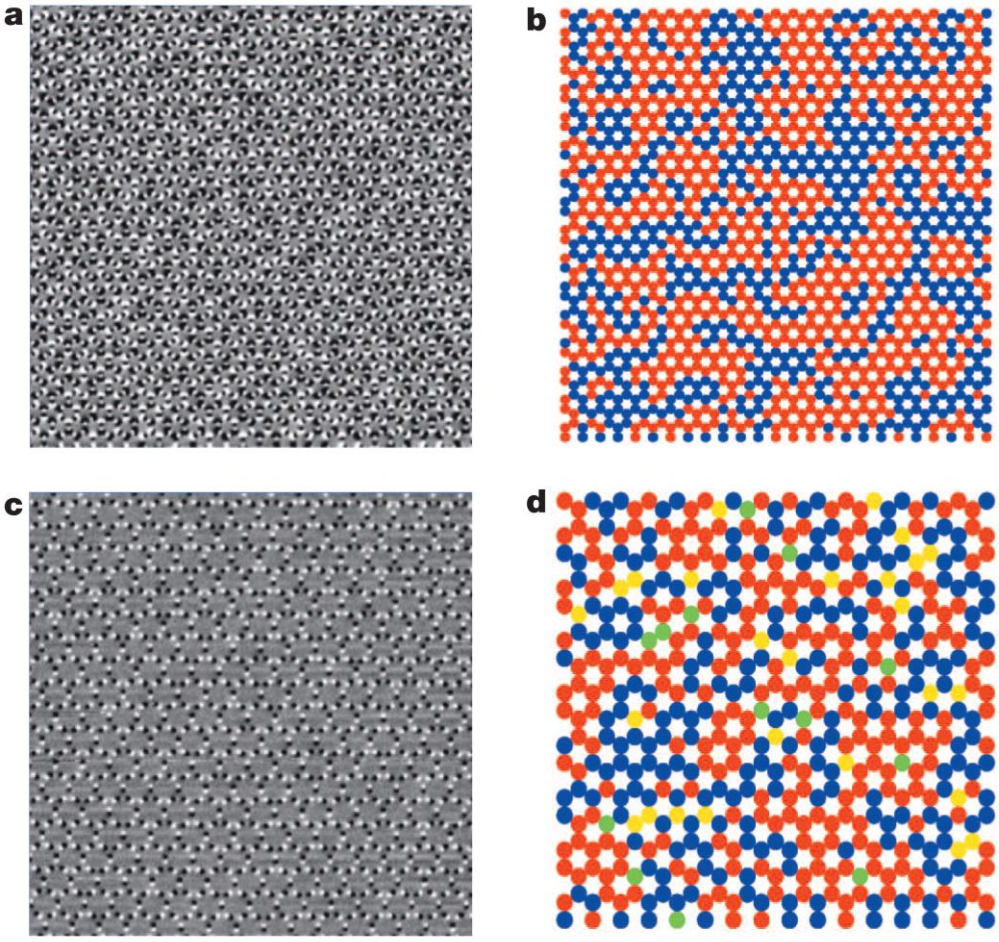

Figure 1.34 **Charge ordering in kagome ice**. **a, c**. MFM maps of thermalized kagome ASI crystals made from permalloy with lattice constant (**a**) 260 nm and (**c**) 490 nm. **b, d**. Corresponding charge domain maps showing incipient regions of charge-ordering. Blue and red dots correspond to vertices within the two degenerate charge-ordered states while green and yellow dots represent ice-rule violating vertices. The smaller lattice constant crystal shows better ordering because it has stronger magnetostatic interactions between neighboring elements. Figure adapted from Ref. [78]**.**

1.4.4 Other lattices

In addition to square and kagome 2D arrays fabricated from magnetic materials, there are a variety of other materials-by-design, artificially frustrated systems that have been studied both experimentally and theoretically. The shakti and tetris lattices are derivatives of the square ASI lattice and are fabricated from nanomagnets, but with selected elements intentionally removed [79,85,147,148]. These systems become more complex as they have vertices where 2, 3, and 4 elements meet. Figure 1.35 shows the structure of the two lattices. The 4-element vertices have the same energy landscape as in square ASI samples, but the 2- and 3-element vertices are new and different. The ground state of these geometries typically has all of the 4-element vertices as type I and all of the 2-element vertices in head-to-tail configurations (type a or type α in Figure 1.35(c)).

The frustration in the system and degeneracy of the ground state arises from the assignment of the 3-element vertex types. These vertices have a lowest energy configuration, type A in Figure 1.35(c), but it is physically not possible for all of the 3-element vertices to be type A. In the shakti lattice, a maximum of 50% of the 3-element vertices can be type A, forcing type B vertices into the system and creating frustration and disorder from where these vertices are distributed [79]. The system cannot satisfy the lowest energy configuration at each vertex, thus it is frustrated, and there are many configurations with different locations of the type A and B vertices that all have the same ground-state energy. Similarly, the tetris lattice cannot have only type A 3-element vertices and must contain a finite number of type B vertices [85]. The frustration due to these geometric constraints has been termed “emergent vertex frustration” and both of these lattices have been investigated as thermally active systems. An interesting feature in the

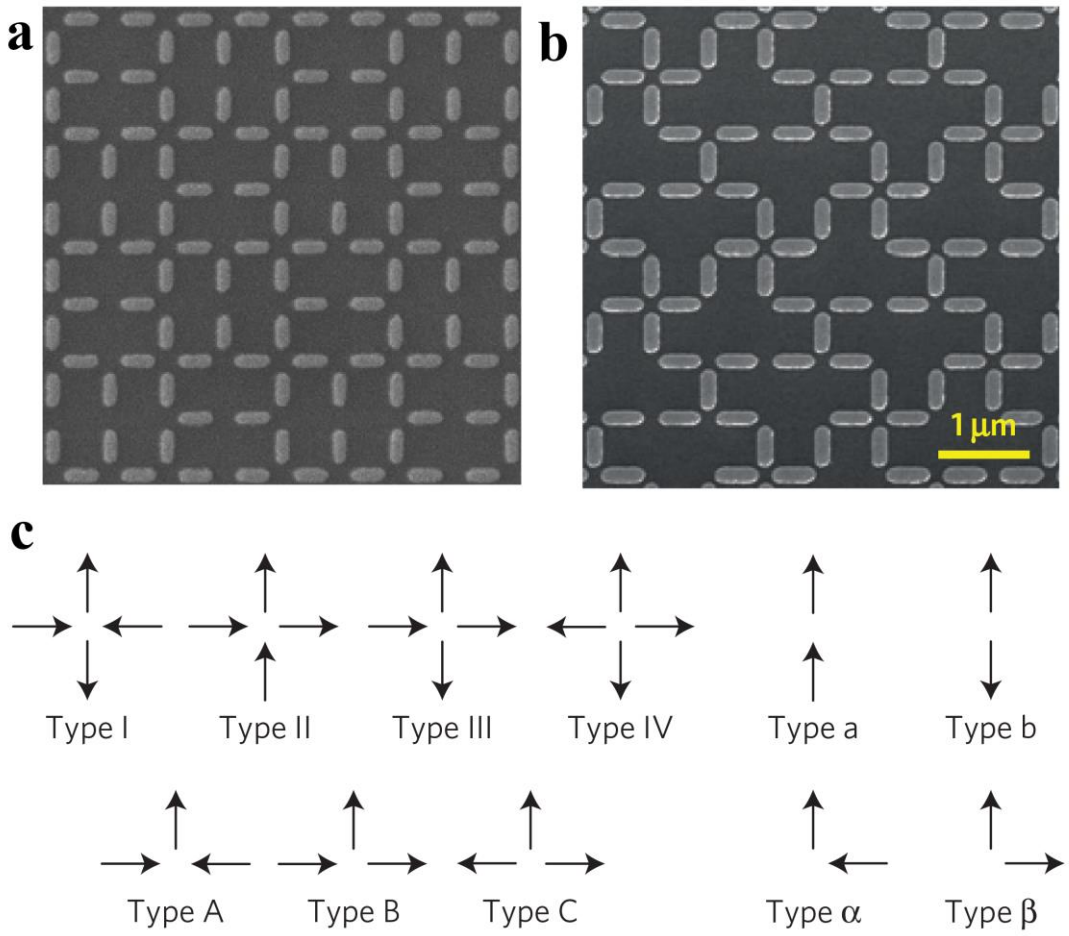

Figure 1.35 **Shakti and tetris lattices**. Scanning electron microscopy (SEM) images of (**a**) the shakti and (**b**) the tetris ASI lattices. Both are derived from the square lattice with some elements removed. **c**. Vertex classifications for these lattices. The type I, II, III, and IV are the same as the square lattice, type A, B and C are 3-moment vertices, and type a and b or α and β are the 2-moment vertices. Figure adapted from Refs. [79] and [85].

shakti results is that there is charge-ordering in the 3-element vertices. The type A and B vertices have a ± *Q* charge on them and these charges are observed to arrange themselves in domains of order when the system is thermalized. Additionally, it is observed that type III 4-element vertices, which have a strong ± 3*Q* charge, attract opposite charges towards them. These charge-charge interactions have a strong theoretical footing, as Monte Carlo

simulations show Coulomb interactions between charged vertices are required in order to reproduce experimental data on vertex populations [79].

Chern and coworkers have proposed a three dimensional ASI system which overcomes some of the issues with the square geometry [149]. 2D square ASI is not an exact analog to pyrochlore spin ice because of the inequivalent interaction strengths between adjacent and opposite spins on the square lattice – one magnet is further away than the other two. This breaks the six-fold degeneracy of the ice ground state. To restore the full degeneracy, Chern *et al*. propose a 3D structure with alternating, stacked planes of nanomagnets. With the correct height between layers, the interaction strength between the four magnets at a vertex becomes equal. This system would be very interesting to study, though it faces some difficult challenges to realize. First, the height between layers must be controlled with extreme precision or else the magnetic interactions would not be equal. Alternating layers need very precise alignment as well. Finally, one must use measurement techniques that can resolve the buried moments in different layers. Improvements in fabrication techniques and development and improvement of 3D imaging techniques could help realize this idea.

Most experimental realizations of ASI have spins with in-plane magnetic moments due to the geometry of the nanomagnet and the chosen magnetic material. In many physical systems, however, spins are not confined to be in a plane. Many models of frustration have out-of-plane spins, notably the 2D Ising model, where Ising spins are perpendicular to a lattice of a given geometry [26]. The picture in Figure 1.13(c) illustrates a 2D Ising model on a triangular lattice with antiferromagnetic coupling. To investigate this and similar models, a version of ASI has been created by Zhang and coworkers using the perpendicular

magnetic anisotropy in Co/Pt multilayer stacks patterned into nanodisks in different geometries [150]. This system has spins pointing out of the 2D plane of the lattice and displays regions of ordering depending on the geometry and coupling strength between the nanodisks. Other unique realizations of ASI include Penrose quasi-crystals [151,152], which show interesting hysteretic behavior in magnetic reversals, and actual macroscopic bar magnets arranged in a kagome geometry and allowed to rotate out of plane [153].

Finally, frustration has been studied experimentally and theoretically in artificially designed systems without nanomagnets. Vortices in superconductors with lithographically patterned double well pinning sites can be used as up/down spin analogues [130,154,155]. These add an additional element to the system because while normally the sites have one vortex in them, they can have defects with two or zero vortices, producing unphysical behavior analogous to a magnet with two north or two south poles [130,131]. A similar realization of artificially frustrated systems is with interacting colloidal particles in periodic optical traps or soft matter [131,156–161]. These systems can also have doubly occupied or empty traps. Both superconducting and colloidal systems can be thermalized during real temperature changes or by applying external stimulus like current passing through the superconductor [131,158,159].

In this chapter, I have discussed the wide body of work which this dissertation fits into. There are clear challenges I aim to overcome with my work, starting with finding a material system for realizing thermal activation. Improvements in the material and thermalization technique will allow for easier experiments and more flexibility, resulting in more results to compare with theory. Improved experimental methods will enable this work to answer outstanding questions on long-range order in frustrated systems, especially

in the kagome ASI geometry. Additionally, I introduce a new type of frustration, specifically due to topological defects, and show how these affect long-range order in the square geometry, an otherwise ordered and well-understood system.

Chapter 2 Fabrication, experimental, and numerical methods

This chapter will give a complete description of the fabrication process and the experimental and numerical methods used in this dissertation. A major aspect of my research was developing a process to grow thin films of magnetic material with favorable properties for thermally active ASI studies so that I could fabricate high quality ASI arrays from these films. This chapter begins with a description of our material of choice, $FePd_3$, its advantages over other groups' materials, and the process I developed for growing high quality films and characterizing them. I will then give a detailed description of the fabrication process I developed including various tips and tricks that may help future researchers working on this project. Next I will provide an overview and introduction to our most important experimental tool, the transmission electron microscope, specifically describing the Lorentz transmission electron microscopy (TEM) technique we use to characterize our samples and image exact spin configurations. I will describe the heating experiments I performed and highlight the advantages of our samples and experimental techniques over other groups' and show how we overcome previous challenges in the field. In this section I will also present my results on thermally active square ASI as these help demonstrate the advantages of $FePd_3$ as well as provide concrete examples as I introduce the experimental and numerical techniques I use. This chapter concludes with the theory and details of the Monte Carlo simulation techniques I employ as well as a framework for modeling the various different experiments discussed in this dissertation. In addition, I will present the results of my Monte Carlo simulations modeling the thermally active square

ASI geometry. Overall, this chapter will detail the methods used throughout this research and the remaining chapters of this dissertation.

2.1 $FePd_3$ as a material for studying thermally active ASI systems

2.1.1 Introduction

An ideal material for thermal activation of ASI would exhibit simultaneously a high magnetic moment per atom and a relatively low T_C. The high moment allows for strong magnetostatic interactions between the macro-spin bar magnets. On the other hand, the lower T_C – ideally not far above room temperature – allows accessing thermal activation while avoiding high temperatures where nanostructures are prone to degradation. Unlike previous studies, which have utilized thin magnetic layers with weak interactions [81,82,84,85,87] or thick layers and high temperatures [78,79], in this chapter I report my work developing a materials system for ASI studies that simultaneously incorporates low temperatures, thick layers, and strong magnetostatic interactions, which are all highly desirable features for thermally active ASI.

I developed techniques for a FePd alloy, close to the stoichiometry of $FePd_3$, for my thermally active ASI samples. To the best of my knowledge, $FePd_3$ has not previously been used for such studies. According to binary alloy work [162], the Fe_xPd_{1-x} system has a local maximum in the T_C, centered on the $FePd_3$ composition, so that fabrications of films within a wide range near this stoichiometric ratio, about 5%, still result in uniform properties and high experimental utility. Nevertheless, $FePd_3$ is a well-defined crystal structure, with a wide solid-solution concentration window. The phase diagram for the Fe-

Pd system is given in Figure 2.1 with the experimentally measured T_C as a function of composition highlighted. At the $FePd_3$ stoichiometry, the T_C of the bulk material is reported as ≈ 260 °C [162–164]. The bulk material can also have a moment density as high as 650 kA m^{-1} [165] which is comparable to permalloy and other strongly magnetic materials. The T_C is ideal for ASI studies because the moments are stable at room temperature, but can be heated easily and quickly, not too hot and far from where the nanomagnets will begin to structurally degrade or melt, in order to see spontaneous flipping of the magnetization of individual magnets.

The properties of $FePd_3$ have some dependence of how well ordered the atoms are in the lattice. $FePd_3$ is known to form a face centered cubic (FCC) $L1_2$ ordered phase, shown in Figure 2.2. This structure has the Pd atoms on the faces of the cube and the Fe atoms on the corners. A perfectly ordered alloy will have all its atoms in these locations, though this is generally never the case and there is always some level of randomness and atomic disorder in the system. The degree of disorder can affect the material's properties because the different atomic magnetic moments and exchange interaction between the atoms depends on the exact interatomic distance [164]. As we will see, annealing films of $FePd_3$ to allow the atoms to find more ordered positions greatly improves the magnetic properties of this material for ASI studies.

The FePd system has unique magnetic properties due to the interaction between the Fe and Pd atoms. Pd is not intrinsically magnetic, but coupling with the Fe atoms can induce a magnetic moment. This occurs through a short range exchange interaction between the Pd 4d and the Fe 3d electronic orbitals which results in a hybridization of these

orbitals [164,168,169]. This splits the Pd 4d band into two bands, one for each spin orientation. The density of states for each of these bands is different, causing an imbalance

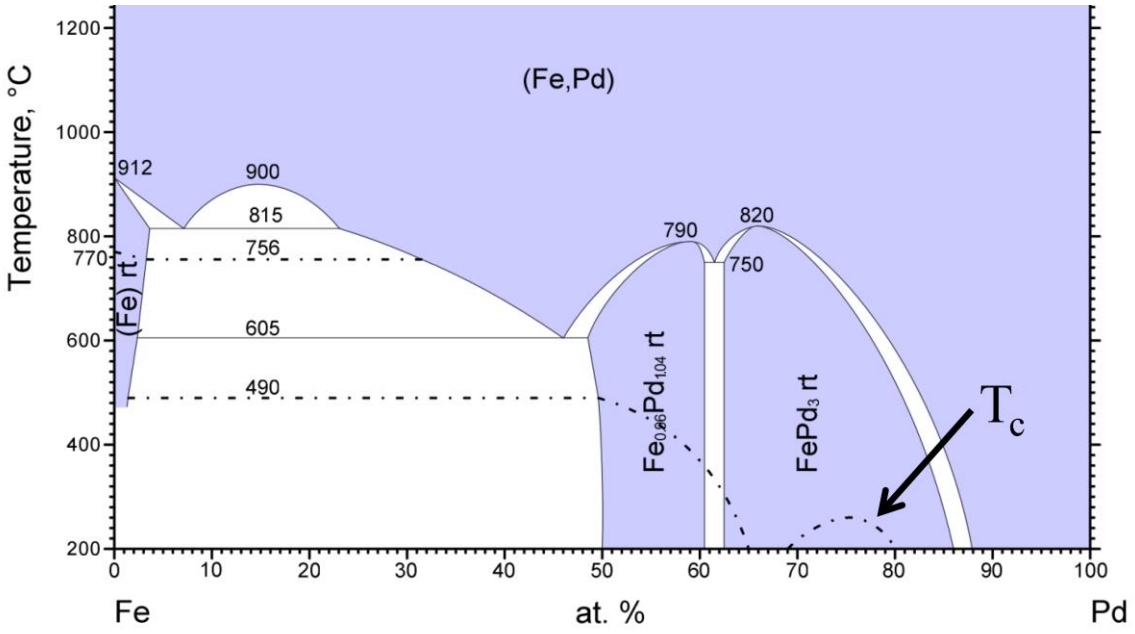

Figure 2.1 **Phase diagram of the Fe-Pd system**. Temperature and composition dependent binary phase diagram. T_C is given by the dotted line where there is a local maximum centered on the $FePd_3$ composition. Figure adapted from Ref. [166].

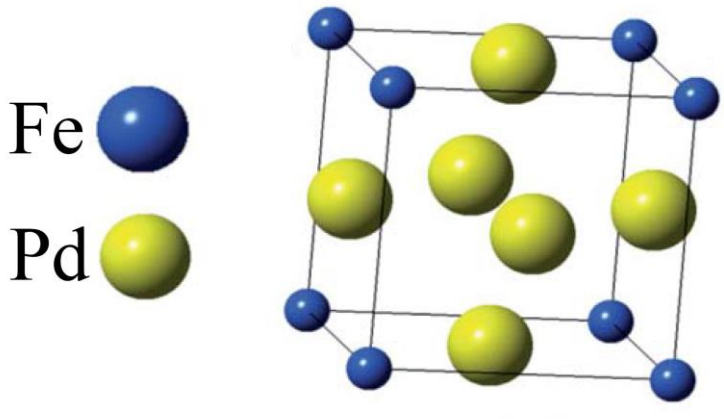

Figure 2.2 **$L1_2$ ordered phase**. Crystal structure of ordered $FePd_3$. The lattice is FCC with Fe on the corners and Pd on the faces of the cube. Figure adapted from Ref. [167].

in the two spin populations and an induced net magnetic moment of the Pd atom [164,170]. The hybridization also affects the Fe moments, with an increasing atomic Fe moment for higher Pd compositions. At the $FePd_3$ composition, the Fe moment is relatively high due to the high concentration of Pd in the system. Elemental Fe has an atomic magnetic moment around 2.2 Bohr magnetons (μ_B) [17]. The calculated and measured Fe moments in $FePd_3$ agree well with each other and are approximately 3.1 μ_B, considerably higher than elemental Fe due to the interaction with the Pd. The Pd moment is ~0.4 μ_B, giving a total moment per $FePd_3$ fundamental unit of 4.3 μ_B, or about 1 μ_B per atom [164,171,172]. Additionally, the unique exchange interaction between the Pd 4d and the Fe 3d orbitals is also responsible for the relatively low T_C of the material. The short range exchange between the orbitals has a relatively small energy difference between the ferromagnetic and paramagnetic states which results in the low T_C.

The value of the magnetic moment for each atom varies with composition but also with the lattice constant of the material, or equivalently the spacing between neighboring atoms. Burzo and Vlaic report a calculated spontaneous onset of the large Fe moment and modest Pd moment when the lattice spacing reaches 3.4 Å and a maximum calculated moment for each atom at 4.3 Å [164]. Under 3.4 Å the exchange interaction causes the atoms to have no magnetic moment as there is a transition to antiferromagnetic exchange below this critical value [164]. The equilibrium lattice constant for $FePd_3$, both calculated and experimentally measured, is approximately 3.85 Å [164,172]. This explains why annealing the $FePd_3$ films is essential. If the atoms do not find highly ordered states with equilibrium lattice spacings or interatomic distances greater than 3.4 Å, the magnetic moment of the atoms are significantly reduced. Annealing helps the atoms move into

equilibrium locations where they are spaced correctly to have favorable exchange interactions and strong magnetic moments.

$FePd_3$ has three other properties that make it favorable for thermally active ASI systems. First, the system has very low magnetocrystalline anisotropy. Strong crystalline anisotropy can compete with the shape anisotropy of the elongated bar magnets in ASI and reduce the success of the Ising-like behavior. A low crystalline anisotropy also makes it easier to fully reverse the nanomagnet and result in all magnetic grains orienting along the long axis, increasing the total moment of the bar magnet. $FePd_3$ is known to have a first anisotropy constant K_1 = -2.0 x 10^3 erg cm^{-3}, which is very low compared to, for example, cobalt, which has K_1 = 4.1 x 10^6 erg cm^{-3} [17,165]. Second, the material has a low coercivity, so that the spins are able to flip easily under the influence of nearby magnets when in a thermally active regime. The low anisotropy and low coercivity make $FePd_3$ what is known as a "soft" magnet. Third, the magnetization of many materials near T_C is known to follow a behavior given by [21]

$$M(T) \sim (T_C - T)^{\beta} \tag{2.1}$$

where β is a material dependent parameter. β has been measured for $FePd_3$ and is found to be ~0.37 [163]. This is favorable because the magnetization will decrease slowly as T approaches T_C, so there is a wide temperature range where the magnets will be thermally active and still posses a finite magnetization. If the transition at T_C were very sharp, it would be difficult to find temperatures that allowed for thermal activation and magnetic relaxation in the lattice. In the following sections, I will present measurements of my $FePd_3$ films that confirm all of the favorable magnetic properties discussed in this section.

2.1.2 Film growth

The ultimate quality of the data from thermally active ASI samples depends heavily on the quality of the films they are fabricated from. Thus, growing high quality films of $FePd_3$ is essential to success of this project. This section will give a detailed description of the film growth process I developed for this project. It is somewhat specific to the tools and resources available at the University of Maryland (UMD) and to the experimental characterization techniques I use as well, though researchers from other institutions should still find the information useful.

I use TEM as my main experimental tool and as I will discuss in Section 2.3, TEM has many advantages over other typical ASI characterization techniques. In order to effectively use TEM to image my samples and their precise magnetic configuration, samples must be fabricated on a substrate that is electron transparent – that is, a substrate that does not have contrast of its own. The Lorentz TEM technique and associated image processing required to characterize exact spin directions is extremely sensitive to small differences in contrast in the TEM images, thus it is essential that the substrate does not add any additional contrast.

The substrate of choice for my work is a SiN membrane, close to the stoichiometric ratio Si_3N_4. A Scanning Electron Microscopy (SEM) image of one is shown in Figure 2.3. The substrate consists of a 3 mm square Si frame, 200 μm thick with a 500 μm square window cut out of the middle. On top of the whole frame is a 100 nm thick, electron transparent SiN membrane. ASI samples are fabricated on the membrane suspended over the cut-out window. While the substrates are necessary for TEM studies, there are some challenges in working with them. The 100 nm thick membranes are very fragile and can

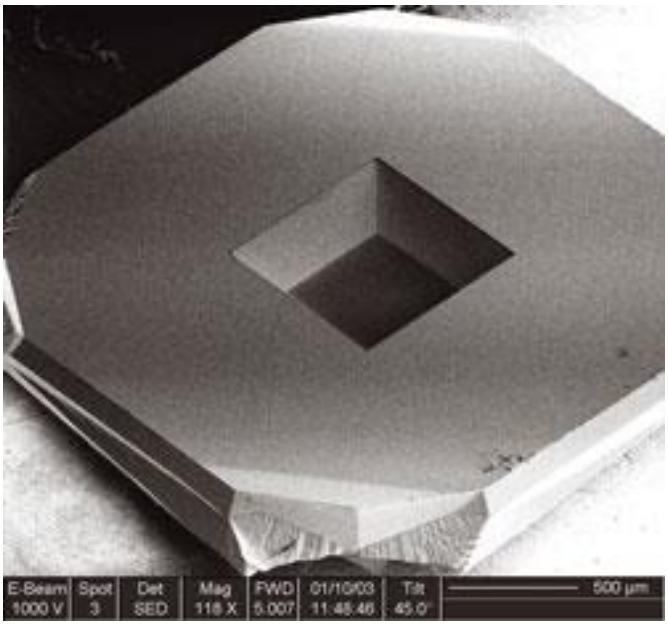

Figure 2.3 **SiN membrane substrate**. Scanning electron microscopy image of a SiN substrate for TEM studies. The substrate consists of a 200 µm thick, 3 mm square Si frame with a 500 µm window cut out of it. On top of the whole Si frame is a 100 nm thick SiN electron-transparent membrane. Figure adapted from Ref. [173].

break easily if dropped or touched. Also, samples must be fabricated one at a time or in small batches compared to other Si wafer based processing. Still, the advantages of using TEM outweigh the challenges of working with the SiN substrates.

Before the $FePd_3$ is deposited, the substrates are prepared with thin underlayers. First, a 2 nm Ti adhesion layer is electron beam evaporated onto the bare SiN membrane. Electron beam evaporation will be discussed further in Section 2.2.2. This layer is important because Ti bonds strongly to the SiN and prevents the films from delaminating as Pd adhesion to SiN is not always very strong. Second, a 2 nm Pd wetting layer is also electron beam evaporated on top of the Ti. This is the layer the $FePd_3$ will be deposited onto and helps facilitate diffusion of the $FePd_3$ during annealing, which is essential to achieving a high quality film. The wetting layer helps control the microstructure of the film, reducing surface roughness and nucleating small grain sizes in the $FePd_3$, both

important factors in the film quality. Small grains help keep the magnet soft and reduce the effects of any magnetocrystalline anisotropy in the material.

After preparing the substrates, they are removed from vacuum and transferred to a sputtering system where the $FePd_3$ is deposited. The sputtering process is shown schematically in Figure 2.4. The technique uses a vacuum chamber filled with Ar gas and a gun with a target of the desired deposition material mounted on it. A large RF potential is applied to gun relative the chamber which ionizes the Ar atoms in the gas to a +1 state, creating a plasma. The plasma is confined to an area near the target using strong magnetic fields from the gun and the charged Ar^+ ions are accelerated towards the target with an electric potential. When the Ar^+ ions hit the target, they knock atoms of the target material off the surface which then diffuse through the Ar gas down to the substrate mounted below (see Figure 2.4).

To deposit $FePd_3$ alloy films I made a custom sputtering target from a pure Fe target decorated with Pd foil. A photograph of my target is shown in Figure 2.5. Essentially, the ratio of the surface area of the exposed Fe to the exposed Pd dictates the composition of the films deposited. I made the target modular so that I could tune the composition by adding or removing pieces of Pd foil until the films were near the desired stoichiometry. To fabricate the target, I first spot welded pure Fe wire in two concentric octagons. I was then able to slide pieces of Pd foil under the Fe wire to secure them to the target. In the next section, I will discuss how I characterized the composition of the films to know when I had reached the $FePd_3$ composition.

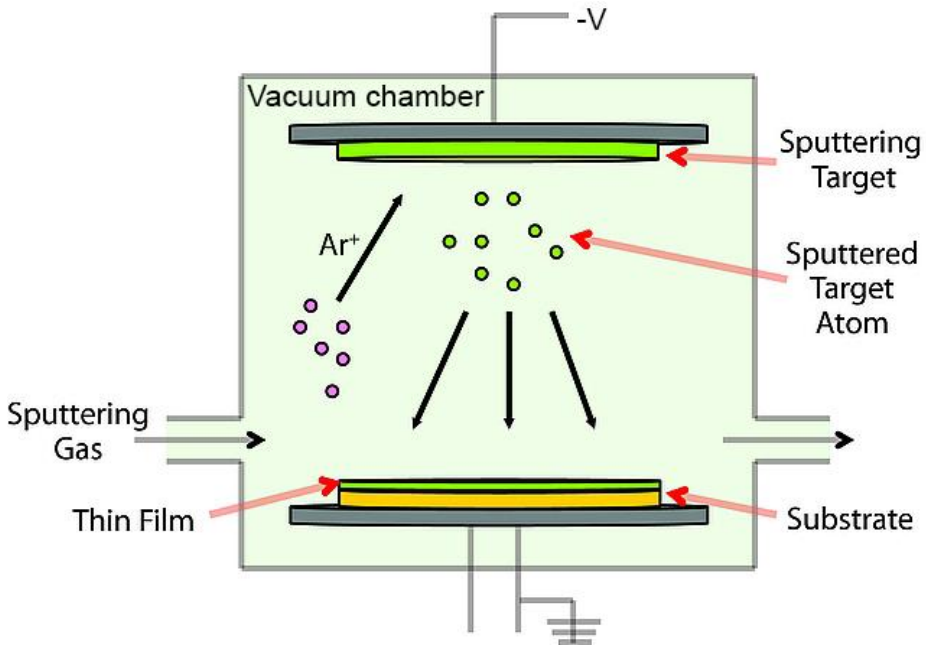

Figure 2.4 **Sputtering**. Gas is fed into a vacuum chamber which is then ionized, creating a plasma, and accelerated at a source target. The sputtering gas ions knock material off the target which diffuses back through the gas and lands on a substrate, growing a thin film of the source target material. Figure adapted from Ref. [174].

The first step in my sputtering process is to load the samples into the vacuum chamber in the Center for Nanophysics and Advanced Materials (CNAM) at UMD and pump it down to base pressure. I would typically load 6-9 individual membranes with the Ti/Pd underlayers into the chamber for each deposition. I often also loaded small Si wafer dies as witness samples for the deposition. The membranes and dies were placed on a small Al holder which was then loaded onto the stage of the sputter chamber. I pump down the chamber first with a rough mechanical pump and then cross it over to a diffusion pump. The base pressure of the chamber varied over the years I used it, depending on what materials other users were depositing, and ranged from 3.0×10^{-6} Torr to 1.1×10^{-5} Torr. Though I always tried to get the base pressure as low as possible, I saw no correlation between a given pressure and the quality of the films grown at that pressure. After the chamber reaches base pressure, I fill it with flowing Ar gas. I adjust the flow rate and pump

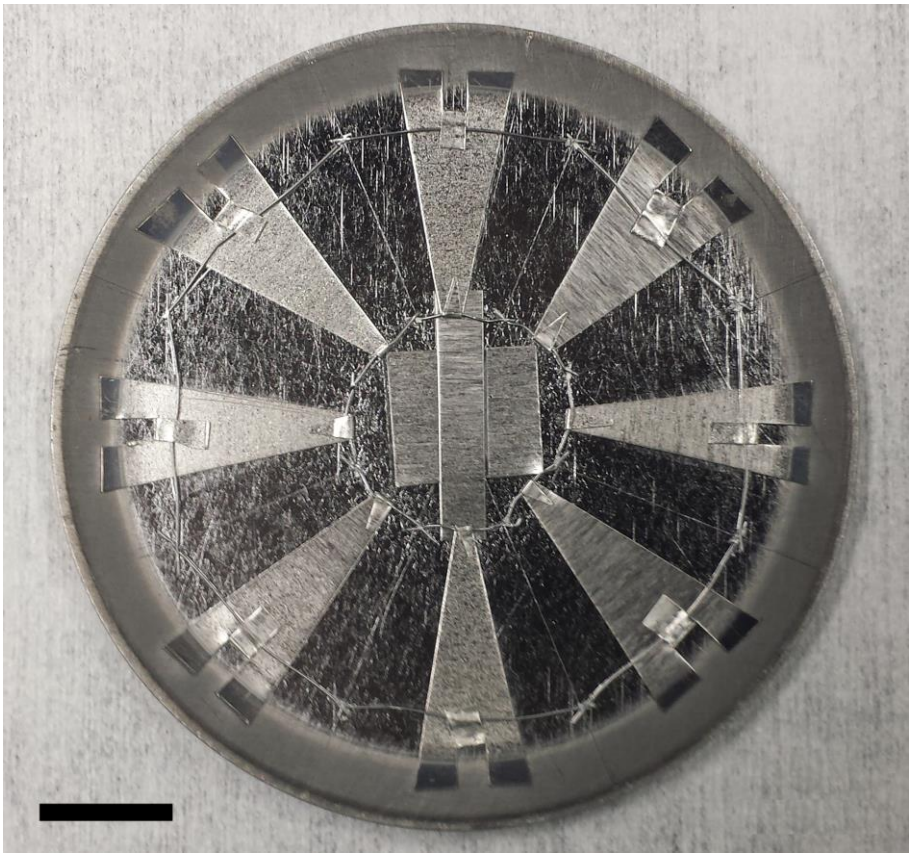

Figure 2.5 **$FePd_3$ sputtering target**. My custom built target is made from a 3 inch pure Fe target with Fe wire welded onto it. Pieces of Pd foil are slid under the wire. The ratio of the surface area of the Fe to the Pd dictates the composition of the film that is deposited. Scale bar is 0.5 inches.

rate of the diffusion pump (by closing the gate valve of the diffusion pump most of the way) to get an Ar pressure of 4.0-4.3 x 10^{-3} Torr in the chamber. I then spark the plasma with 200 W RF power. There is a shutter that protects the membranes while the system pre-sputters for a few minutes to remove any oxide layer that has formed on the target since the last sputtering run. The color of the plasma changes from pink to purple once the oxide layer is removed. I then open the shutter and deposit 23 nm material onto the membranes at a rate of about 1.1 Å s^{-1}. The deposition rate, as well as the composition of the films

deposited, depends on the sputtering power. I generally kept the power constant at 200 W for consistency.

The deposition rate and total thickness are measured using a quartz crystal and crystal monitoring system which I calibrated for my target and other deposition parameters. I used the density and Z ratio for Pd and changed the tooling factor to account for the alloy composition. To calibrate, I would bake drops of resist (see Section 2.2.1) onto Si witness samples and deposit 100-150 nm of material onto the samples. I would then remove the drops of resist and measure the step height using a profilometer. Averaging over several steps on one sample and over several different samples generally gave an accurate measurement of the actual film thickness. I would then use the experimentally measured thickness compared to the crystal monitor reading to adjust the tooling factor and would repeat the whole procedure if necessary.

As-deposited films tended to be very rough, columnar, and generally not even ferromagnetic. An annealing procedure was essential to making the films magnetic, reducing the surface roughness, and achieving a nice grain structure. Figure 2.6 shows before and after annealing images of the films. The sputter chamber I used to deposit the films did not have a heating stage, so I would remove the films from that chamber and load them into a different system in Professor Ichiro Takeuchi's lab at UMD. I placed them on a steel holder and then positioned the holder on a heating stage in an Ultra-High Vacuum (UHV) chamber. I would anneal the films at 750 °C for three hours in UHV (base pressure $< 6\times10^{-8}$ Torr). There are numerous papers that report an improvement in $FePd_3$ magnetic properties after annealing at 600 °C or higher [162–165]. After annealing, I would

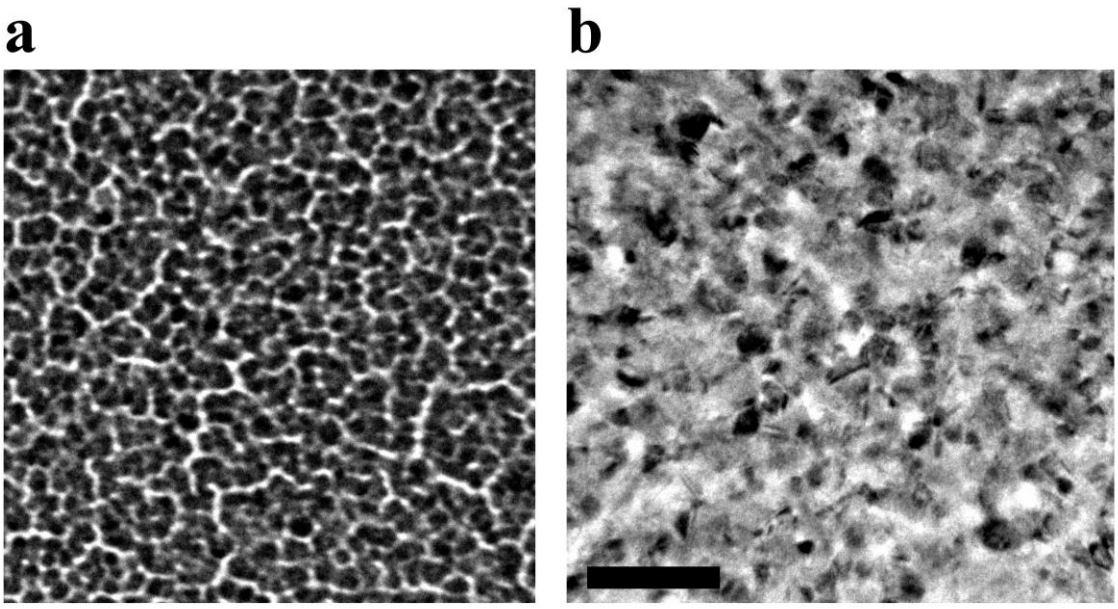

Figure 2.6 **Annealing of $FePd_3$ film**. High resolution TEM images of $FePd_3$ films (**a**) before and (**b**) after annealing showing the grain structure of the film. Before annealing, the film has a columnar structure with large gaps and cracks. After annealing, the film is continuous and grains are 3-7 nm in diameter. Scale bar is 30 nm.

characterize the films for their composition and magnetic properties. The characterization techniques I used are described in the next section.

2.1.3 Film characterization

Once $FePd_3$ films are fabricated, it is essential to check their properties before patterning them into ASI arrays. The main characteristics I checked were the composition, the grain structure, the magnetization and how well they supported magnetic domains and domain walls, and their T_C. Much of this characterization can be done with TEM, though I used other tools as well. During the initial stages of developing the film growth process I was only measuring the composition of the films in order to tweak the design of my alloy target to get the composition near $FePd_3$.

Composition is characterized using Wavelength Dispersive x-ray Spectroscopy (WDS) on a JEOL JXA-8900 electron probe microanalyzer. WDS is a very quantitative and precise technique for measuring the composition of materials. A beam of electrons is accelerated towards the material at high energy which knocks out core electrons of the atoms. This leaves the atoms in a very unstable state and an outer shell electron quickly transitions down to fill the core shell vacancy, emitting an x-ray with an energy characteristic of the specific element it came from. The x-rays are collected and used to identify the elements present and their relative concentrations. Often, the x-ray energies are measured directly in a technique called Energy Dispersive x-ray Spectroscopy (EDS). However, this technique does not have very precise energy resolution and it can be difficult to distinguish different elements that have characteristic x-ray energies very close to each other.

Since the wavelength of a photon is inversely proportional to its energy by hc, where h is Plank's constant and c is the speed of light, the x-ray can be equally well characterized by its wavelength. WDS tends to be more precise than EDS in quantitative measurements of composition because it can better distinguish wavelengths that are close together. The x-rays are passed through different crystals which diffract them using Bragg's law [17]

$$2dsin\theta = n\lambda \tag{2.2}$$

A detector in the microprobe is moved to different θ to help detect the different characteristic wavelengths. The microprobe analyzer also has a built-in computer with calculated and well characterized x-ray wavelengths and emission rates for different elements as well as sophisticated algorithms for determining relative emission rates for the

different characteristic x-rays. Before each measurement, these rates and the exact elemental wavelengths are re-calibrated from samples of pure Fe and Pd always kept inside the tool and this aids is giving highly quantitative and precise elemental information. I performed numerous depositions from my FePd alloy target, adding additional pieces of Pd foil until I was consistently getting compositions within a few atomic percent of the $FePd_3$ stoichiometry.

After I finalized my target design, the exact composition of my films still varied slightly from run to run due to the chaotic nature of sputtering and my nonstandard target design. However, in a single deposition there was generally very uniform composition, only varying by 2-3% across several substrates spread out over an 8 cm deposition area, with each 3 mm substrate being highly uniform in composition. Different runs gave slightly different magnetic properties and sometimes different T_Cs dependent on the given composition. For each batch of films, I would measure the T_C using Lorentz TEM and a specialized heating holder. A schematic of the heating holder is shown in Figure 2.31 and this holder will be discussed further in Section 2.4.

Figure 2.7(a) shows a Lorentz TEM image of one deposition of a 23 nm thick $FePd_3$ film at room temperature. The textured bright and dark contrast in the image indicates the magnetization of the film is relatively high as the amount of contrast is proportional to the magnetization in the film [175]. Figure 2.7(b) shows that when the film is heated to 140 °C, the Lorentz contrast disappears, indicating the material is near or above its Curie temperature. Some early films I made had a T_C closer to 110 °C, but most of the later ones had a higher T_C, around 140 °C. I will note that measuring the T_C using this method is not perfectly accurate due to some considerations with the *in-situ* heating holder I will discuss

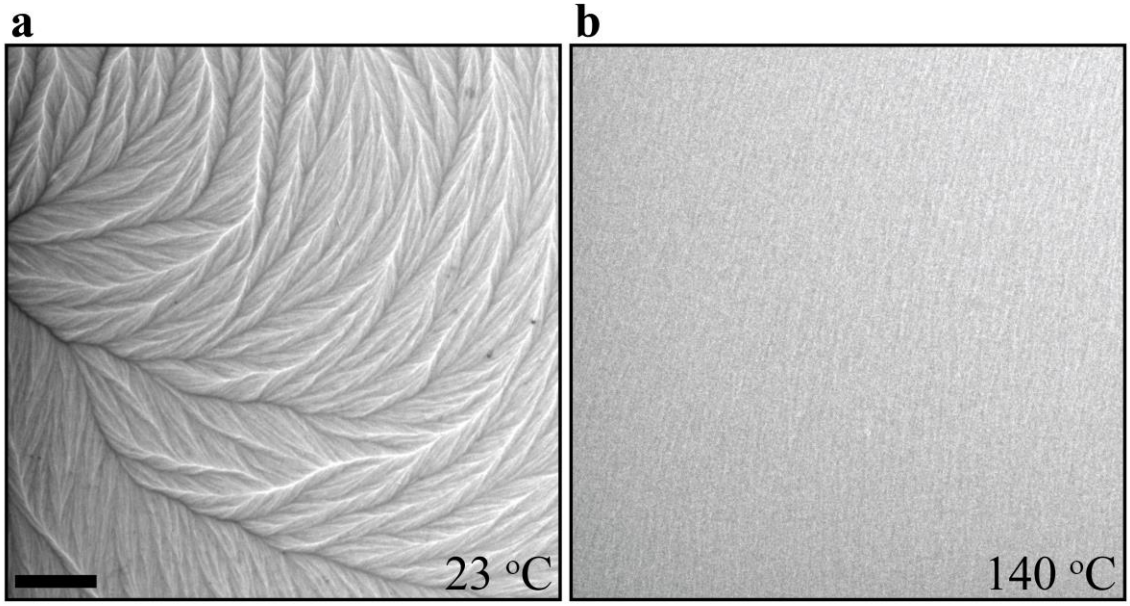

Figure 2.7 ***T_C* characterization**. Lorentz contrast TEM images of a 23 nm $FePd_3$ thin film at (**a**) room temperature and (**b**) 140 °C. The textured contrast in (**a**) indicates the film has a high magnetization, while the absence of contrast in (**b**) signifies the material is near or above its Curie point. Scale bar is 20 μm.

in Section 2.4. The T_C of the films may have been slightly higher than reported here as I will show with some measurements at the end of this section.

Figure 2.6(b) in the previous section presents a high resolution TEM image showing the grain structure of a film. Grains in my films were typically 3-7 nm in size, which is favorable for thermal ASI studies since it will give rise to a relatively low coercivity. The grain structure is generally very consistent between different depositions and is a good way to check that the annealing process was successful. One of the largest variations between different film growths is the magnetization of the films and how well they support magnetic domains. A film with nicer domains and sharper domain walls in the Lorentz TEM images tended to have better Lorentz contrast in the individual patterned nanoscale bar magnets

and yielded better data and more-ordered states than a film that only displayed ripple Lorentz contrast.

An example of the two main types of contrast observed is shown in Figure 2.8. The films with sharper domain walls have a higher magnetization which is better for ASI studies and probably have less surface roughness as well. The variability of the total magnetization and surface roughness could depend on many factors, though most likely it is the sputtering process and the annealing of the films that have the greatest effect. Some batches of films came out better than others and I did not have as much control as I wanted over the quality and consistency of them. I used the best films with the best magnetization and sharpest domain walls for my samples and the data in this dissertation, but I made many other films that did not turn out as well as those. At the end of this section, I discuss some possible improvements in future work to overcome these challenges and increase the consistency of the $FePd_3$ films.

Finally, a group at the University of Illinois that also studies thermal ASI samples was generous enough to take some SQUID microscope measurements of my $FePd_3$ films. Figures 2.9 and 2.10 confirm all of the favorable magnetic properties of my films discussed in this section and are credit to Ian Gilbert and Peter Schiffer at the University of Illinois. The only discrepancy between what I characterized and their measurements is the T_C of the film which appears to be slightly higher than 140 °C, the temperature where I saw the Lorentz contrast disappear in the TEM images. This may indicate that there is still some magnetization in the film when the Lorentz contrast disappears, just not enough to give any noticeable contrast. Figure 2.9 shows the magnetization as it approaches T_C, showing a nice gradual approach, though it does not go to zero at 140 °C. The hysteresis curves in

Figure 2.10 confirm that the coercive field is very low and that there is no strong crystalline anisotropy in the films.

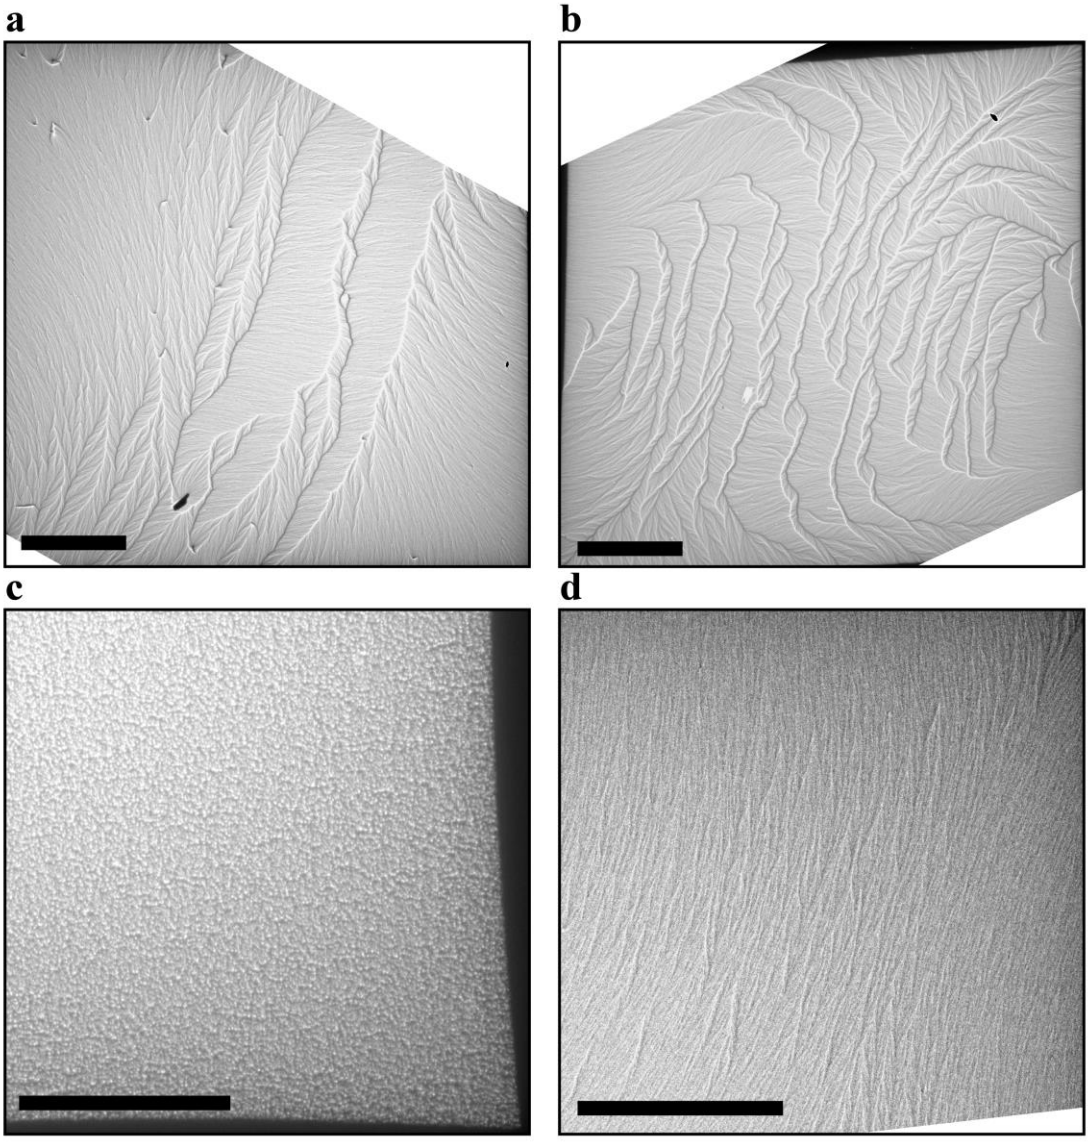

Figure 2.8 **Lorentz contrast in different $FePd_3$ depositions**. **a-d**. Lorentz TEM images of thin films of annealed $FePd_3$. In (**a**) and (**b**), the films have a high magnetization and low surface roughness which supports magnetic domains and shows sharp, strong Lorentz contrast. In (**c**) and (**d**), the films only display ripple contrast, indicative of more surface roughness which is less favorable for fabricating high quality ASI samples. All scale bars are 100 μm.

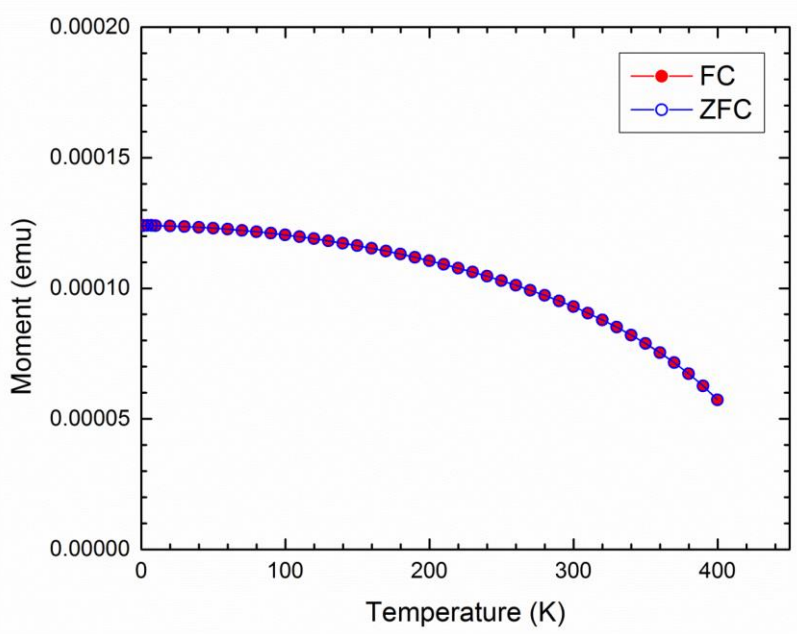

Figure 2.9 **Magnetization as a function of temperature**. Field cooled (FC) and zero field cooled (ZFC) magnetization M vs. temperature T for a thin film of $FePd_3$ showing a gradual approach to T_C. The T_C appears to be slightly higher than was measured in Lorentz TEM heating experiments shown in Figure 2.7, though it is still relatively low. Measurement and figure courtesy of Ian Gilbert and Peter Schiffer at the University of Illinois.

In future work on this project, improving the consistency and quality of the $FePd_3$ films will be essential. I fabricated a number of very good films that yielded useful data, but I also made many films that did not have as desirable properties. I was able to work with the high quality films I had, but it would have been better to be able to readily produce the highest quality films. To accomplish this, it would be beneficial to custom order a uniform $FePd_3$ alloy sputter target to get more consistent depositions. I believe that a careful study of the annealing temperature and time would also help produce the best possible films. I had to break vacuum between depositing and annealing my films and if this process could be accomplished in one chamber it would certainly reduce the time it takes to grow the film and almost certainly improve the properties by reducing the oxidation of the as-deposited film shown in Figure 2.6(a). I did some work to study the

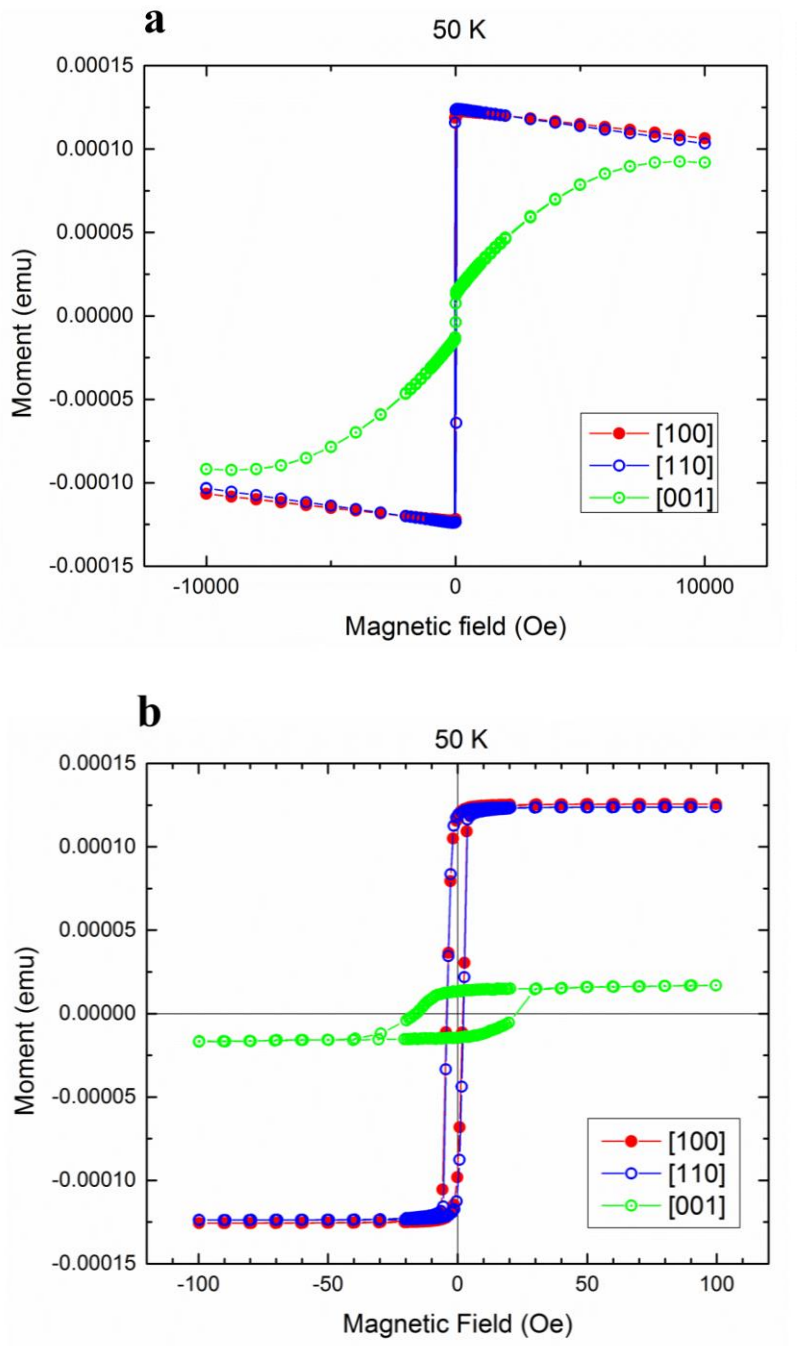

Figure 2.10 **Hysteresis curves**. **a, b.** Magnetization vs. applied field hysteresis curves for a thin film of $FePd_3$. The measurement was taken at 50 K for both in-plane and out-of-plane magnetization. The film shows no sign of in-plane anisotropy which is favorable for ASI studies. (**b**) shows a higher resolution detail of the same measurement in (**a**). Measurement and figure courtesy of Ian Glbert and Peter Schiffer at the University of Illinois.

effects of the underlayers, and tuning the thickness of these or potentially using different material could help as well. Overall, a new target and a fine-tuned annealing procedure all in a single vacuum chamber should improve reliability and yield excellent films with correspondingly good data.

2.2 Fabrication

An important part of my project was to develop a process to pattern my $FePd_3$ films with favorable properties for thermal activation of ASI into usable ASI samples. I built upon Stephen Daunheimer's work in our group making kagome ice samples out of permalloy [132], but had to significantly modify the process to adapt it to the $FePd_3$. I also fabricated some ASI geometries that were different from those made by Daunheimer and altered the process to suit the new geometries. This section will give a broad overview of each technique and tool used in my fabrication process as well as all of the details required to reproduce the process.

An overview of my method is shown in Figure 2.11. Starting with the fully characterized $FePd_3$ films on top of SiN TEM membrane substrates, I spin coat a bilayer polymer resist onto the membranes (Figure 2.11(c), Section 2.2.1). I perform electron beam lithography (Figure 2.11(d), Section 2.2.1) to define the ASI patterns and then electron beam evaporate an Al etch mask over the whole sample (Figure 2.11(e), Section 2.2.2). A liftoff step is performed (Figure 2.11(f), Section 2.2.3) and finally, the pattern is transferred into the $FePd_3$ with Ar ion milling (Figure 2.11(g), Section 2.2.4). Characteristic images of samples I made are shown in Section 2.2.4. In Section 2.2.5, I discuss some

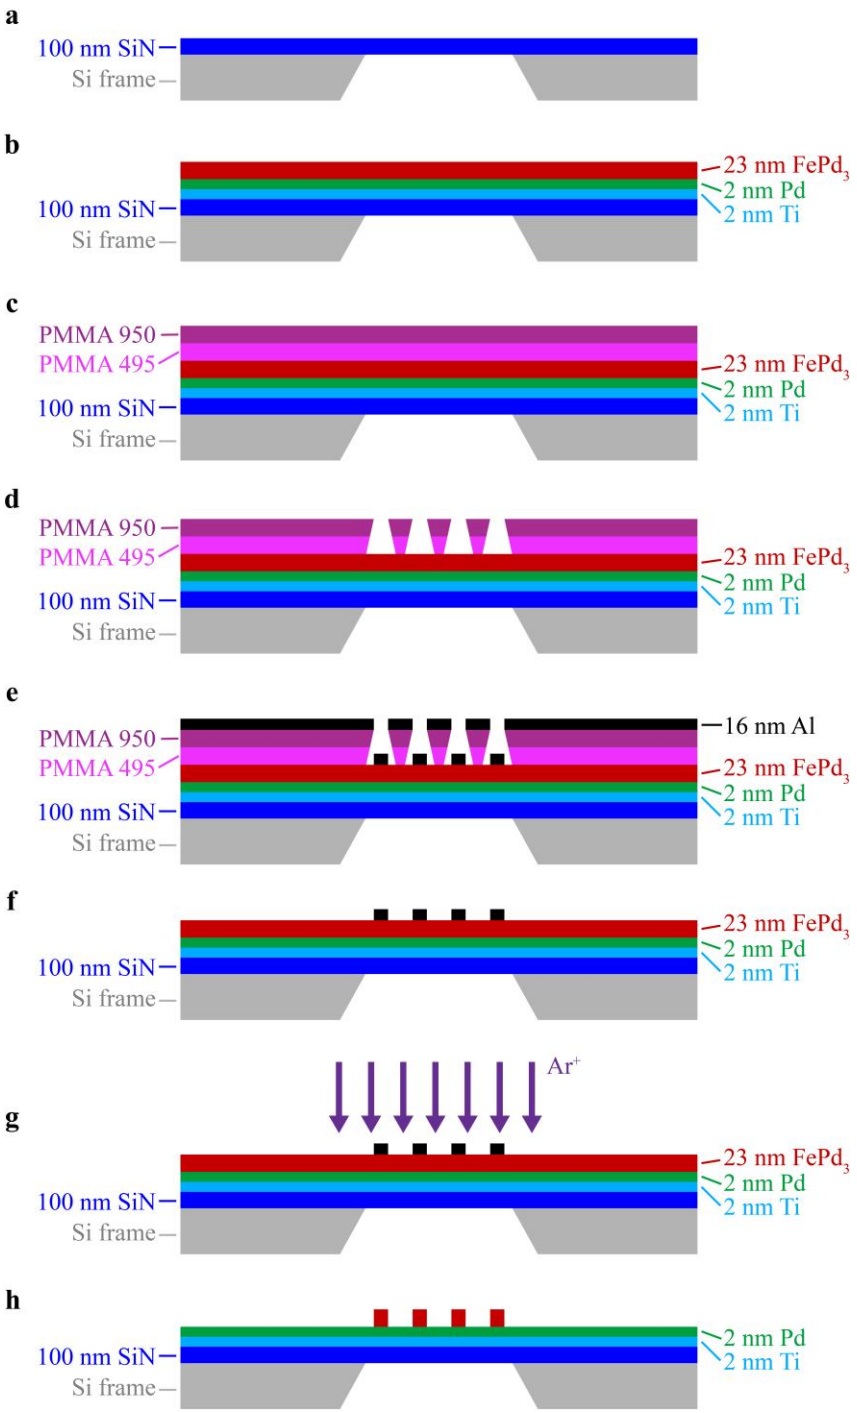

Figure 2.11 **Fabrication process overview**. The sample starts as a bare SiN TEM membrane substrate (**a**). A 2 nm Ti adhesion layer and 2 nm Pd wetting layer are electron beam evaporated onto the membrane. Then a 23 nm $FePd_3$ film is deposited using RF magnetron sputtering (see Section 2.1.2) (**b**). The sample is spin-coated with a bilayer resist (**c**) and electron beam lithography is performed (**d**). A 16 nm Al etch mask is electron beam evaporated over the whole sample (**e**) and liftoff is performed (**f**). The sample is finally etched using Ar^+ ion milling (**g**), transferring the pattern into the $FePd_3$ film (**h**).

common problems and modes of failure for the fabrication process with the intention that these may help future researchers intending to use the procedure I developed.

The method for fabricating samples of $FePd_3$ ASI needed modification from Daunheimer's process [132] for permalloy samples because of fundamental differences in how the magnetic materials are deposited. Permalloy can easily be electron beam evaporated onto lithographically patterned polymer resist and lifted off to define the individual nanomagnets. The problem with using this technique with my $FePd_3$ films is twofold. First, electron beam evaporation is a low pressure process, typically around 3×10^{-6} Torr. The mean free path of atoms at this pressure is around 1 meter, larger than the distance between the evaporation source and sample inside the evaporation chamber (see Section 2.2.2). At this low pressure, the permalloy atoms reach the sample highly anisotropically and will only coat the top of the resist or land on the substrate, similar to the evaporated Al etch mask in Figure 2.11(e). This leaves the resist sidewalls exposed, allowing for a solvent to access the remaining resist to remove it and the metal on top of it during liftoff.

On the other hand, sputtering is a much higher pressure process, around 4×10^{-3} Torr of Ar as discussed in Section 2.1. The mean free path of atoms at this pressure is only on the order of centimeters and their angles of incidence to the sample are nearly isotropic and the atoms end up coating the sidewalls of the resist. This prevents the liftoff solvent from being able to access the remaining resist and will prohibit good liftoff from occurring. In addition to the pressure difference between the deposition methods, the sputter deposited $FePd_3$ needs to be annealed at high temperature to achieve its favorable magnetic properties as was discussed in Section 2.1.2. The continuous thin films of the

material benefit from high heat, but patterned nanostructures will degrade at the annealing temperature due to their small size. Therefore, it is the need to anneal the $FePd_3$ at high temperature and the fact that sputtering is a conformal technique and not compatible with liftoff which motivated the development of a new fabrication process for my ASI studies.

2.2.1 Electron beam lithography

Lithography is a common process for patterning structures and circuits of all kinds used in both research and industrial manufacturing. There are two main types of lithography. The first is optical, or photolithography, which uses visible light passing through a semi-transparent mask to expose regions of a polymer resist sensitive to light. The second is electron beam lithography (EBL), which uses electrons directly traced into a resist sensitive to electrons [176]. Wherever the resist is exposed to light or electrons, the polymers in the chemical react to the stimulus, cross-linking with each other and locally change their structure. After exposure, the resist is developed and depending on the type of resist, the exposed area will be removed and everywhere else the resist will remain or vice-versa, leaving the pattern or its negative defined on the substrate. This is depicted schematically in Figure 2.11(d) for the type of resist I use, known as a positive resist, where the exposed regions are removed.

Photolithography is most common in industrial applications because of its high-throughput capabilities. It is fast and reliable, but requires a static, premade mask specifically designed for a single pattern only. Also, the feature size of the structures that can be made by photolithography is limited by the wavelength of light used, typically

around tens or hundreds of nanometers. There have been many improvements to the process, driven by modern integrated circuit and semiconductor engineering, including a shift towards extreme ultra-violet light which has a shorter wavelength and thus can reduce the minimum feature size. Nevertheless, to get the smallest feature sizes and also produce small features more reliably, EBL is a better method [176]. The de Broglie wavelength of electrons depends on their energy. A typical EBL system accelerates electrons with a potential of tens of kilovolts, giving the electrons a wavelength on the order of 10 pm. This is significantly smaller than anything achievable with optical lithography and can be used to make the smallest possible feature sizes and structures.

Small and reliable feature sizes are essential for ASI studies, which makes EBL the better technique. EBL is also more suited for low-throughput research applications because it is much more flexible. New patterns, sample designs, and geometries can readily be created, tested, and changed. EBL is slower than photolithography, but in ASI research where only a few samples need to be produced at a time, EBL is the better choice. A schematic of an EBL tool is shown in Figure 2.12. Electron-beam writers are typically modified scanning electron microscopes with precise stage control and beam deflectors that quickly and accurately move the beam of electrons, tracing out the path of desired structures.

A major advantage of EBL is that the patterns the tool writes can be readily created and changed. At UMD, I use a Raith e_LiNE system that has an associated proprietary editor for defining the lithography patterns. The patterns are first created with the computer program and then transferred to the system that controls the electron beam and writes them. The Raith system can draw both lines and 2D filled-in shapes. The lines and 2D areas in a

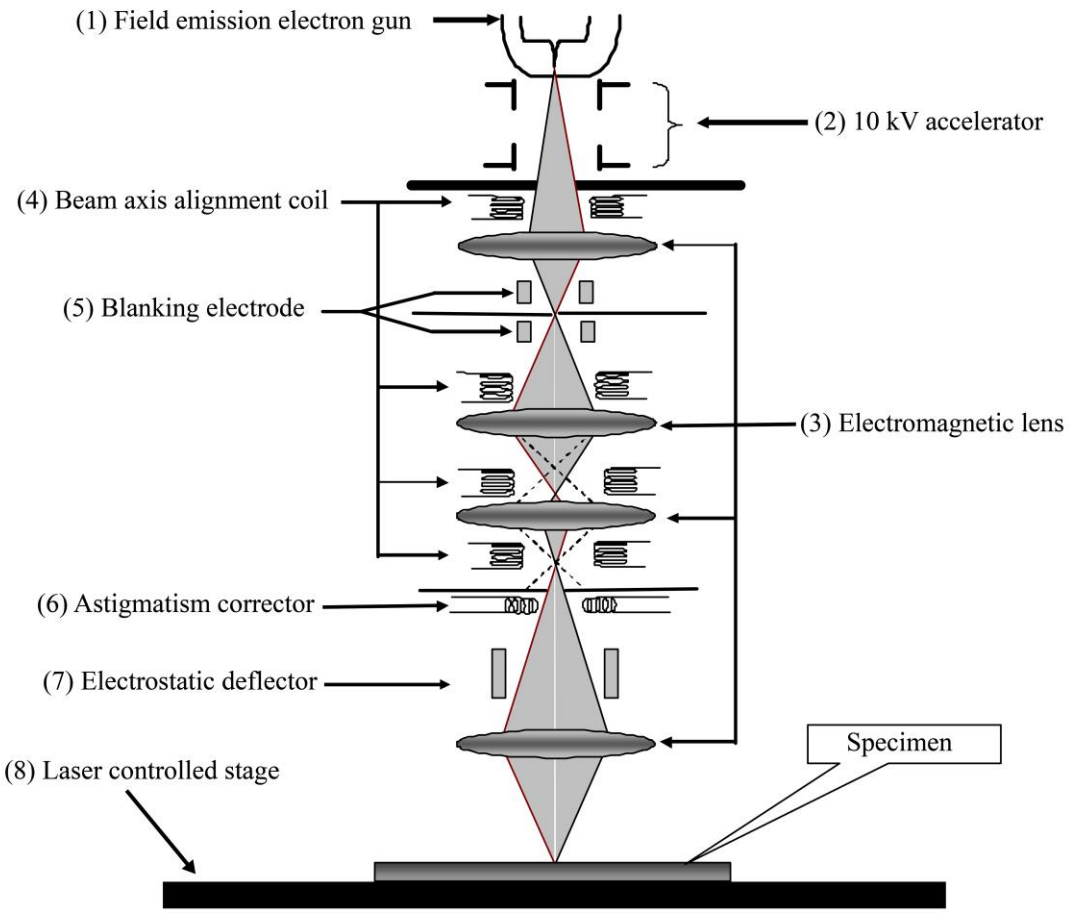

Figure 2.12 **Electron beam lithography**. Lithography capabilities are integrated into existing scanning electron microscopy technology. The electron gun (1), accelerator (2), electromagnetic lenses (3), axis alignment coil (4), astigmatism corrector (6) are all typical components in a scanning electron microscope. The addition of a blanking electrode (5), electrostatic deflectors (7), and a laser controlled stage (8) allow for high speed and high accuracy patterning of structures. Figure adapted from Ref. [177].

pattern have different electron dose settings that can be controlled individually. The dose is essentially how many electrons are accelerated into the resist per unit length or per unit area in order to expose it and remove the resist. There is a certain threshold that must be reached in order to fully expose the resist so it is able to be removed, but too high a dose can cause other unwanted effects. I draw all of my ASI nanomagnets as lines in the lithography software. The length of the line dictates the length of the nanomagnet and the line dose will determine how wide the magnet ends up being, with a higher dose resulting

in a wider magnet. Line doses are carefully calibrated for each geometry to result in magnets of the correct size and these doses are given later in this section.

ASI arrays often have thousands of individual magnets and it would be extremely time consuming and tedious to manually draw every line by hand. Instead, I use a MATLAB script to automatically generate the lithography patterns. Stephen Daunheimer wrote a script to produce patterns for kagome geometry samples. I utilized and modified his script for my kagome samples and I also made new scripts for my square geometry samples and for the topological defect samples discussed in Chapter 4. The MATLAB scripts generate a text file with a particular syntax that can be read by the Raith software. Since the ASI patterns are made entirely of lines, the script will specify to draw a line with a given start and end point, write that line to the text file, calculate the next line's coordinates, and then write that line to the text file, continuing until the whole pattern is written. I then take the text file and import it into the Raith software, checking to make sure it was produced correctly.

A single ASI array of a given size is often referred to as a "crystal." I fabricated kagome crystals as small as one hexagon with six individual magnets or as large as 27 hexagons on a side with over 7,000 elements. Example of these can be found in Chapter 3, for example, in Figures 3.6 and 3.23. The square samples I made were typically 40 elements on a side; an example is shown in Figure 2.16. In my kagome geometry samples, the magnets are connected in order to increase their magnetostatic interactions. In the lithography patterns for these samples, the lines are also connected. Daunheimer figured out that is it helpful to actually write each element twice, going over the resist once in each direction. This also helps to keep the vertices symmetric in the kagome samples which is

important in reducing disorder in the lattice. A schematic of the kagome re-write scheme is found in Figure 2.13(a), showing the symmetry of the write algorithm. The order in which the MATLAB script generates the individual lines dictates the order they are written by the Raith pattering software, thus there is precise control over the pattern and line order which is essential for creating uniform elements and vertices. I found rewriting elements was very important in the square geometry as well to keep the magnets and the vertices symmetric. The scheme for the square rewrite process is shown in Figure 2.13(b).

Many times multiple crystals of different sizes or geometries are fabricated on the same SiN membrane. Sometimes this is done to calibrate the correct electron dose for a new pattern or geometry. To do this, an array of varying dose is generated and after

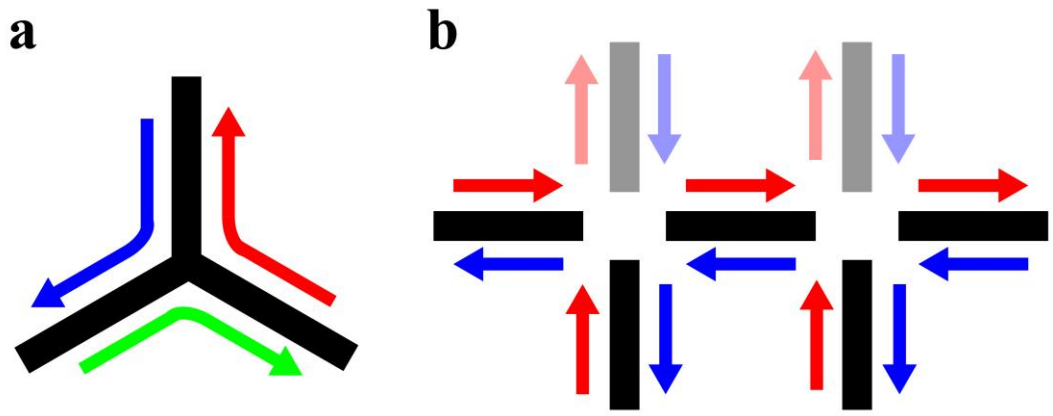

Figure 2.13 **Lithography re-write scheme**. The re-write scheme the EBL pattering software uses to write the (**a**) kagome and (**b**) square geometries. In (**a**), each line is written twice, once in each direction, and each vertex is kept symmetric by moving the electron beam along each of the 120° interior angles. In (**b**), the electron beam moves along one row of elements alternating writing a horizontal and a vertical element. Then the beam moves back along the same row, writing over all of the elements in the other direction. The red arrows represent lines written in the first past and the blue arrows are lines written on the way back. The vertical elements of the next row are also shown in lighter colors.

fabrication, the same is examined to determine what dose yields the desired size elements. Otherwise, fabricating many different crystals on one sample is a good way to test multiple geometries and effectively carry out multiple experiments at once. A great deal of data can be generated from one successful sample with many different crystals on it. Typically, a 100 μm x 100 μm array of different crystals is designed and then this array is repeated in a 4 x 4 grid across the membrane. This provides 16 copies of each crystal or geometry. Many times one part of the membrane will be dirty or have the wrong dose or bad liftoff, but another part will be usable. This technique helps improve the success rate and utility of every sample run.

Once patterns are defined in the Raith software, the first step of EBL is to coat the $FePd_3$ films and substrates with resist (see Figure 2.11(c)). I use two layers of poly(methyl methacrylate) (PMMA) resist with slightly different molecular weights. The bottom layer has a lower weight, which makes it more sensitive to electron exposure, so that the same electron dose will remove a larger area than the top later. This creates an undercut as shown in Figure 2.11(d). The first layer is MicroChem PMMA 495 C2 resist. 495 is the molecular weight and C2 relates to the concentration of the polymer in solution and affects how thick the coating is when spun on at different rates. We use a spin coater in our lab with a special chuck that is specifically designed to hold the SiN TEM membrane substrates. For each layer of resist, I put one drop of the solution into a small groove in the chuck that helps hold the SiN chip in place while spinning and then put two drops directly on top of the membrane. I spin the PMMA 495 layer at 6000 RPM for 45 seconds, with a 5 second acceleration and 5 second deceleration. I estimate the thickness of the resist spun at this rate to be around 100 nm. The resist is then baked at 180 °C in an oven for 45 minutes.

Baking the resist removes the solvent that the polymer is suspended in and leaves just the polymer on the samples. The next layer is MicroChem PMMA 950 C2 resist, again spun at 6000 RPM for 45 seconds with a 5 second acceleration and deceleration. The second layer is also baked at 180 °C for 45 minutes. Finally, a layer of Mitsubishi Rayon's aquaSAVE conductive polymer is spin coated in the same way as the PMMA. This layer is not baked and is not a resist. Its purpose is to prevent the PMMA from chargeing up with electrons during lithography, as this would deflect the electron beam or make it more difficult to focus the beam. The aquaSAVE prevents charging by giving the electrons a path to move off of the resist surface.

Spinning resist is one of the most variable steps in my fabrication process. Because of the SiN substrates' small size, the resist often does not coat evenly or reproducibly. The thin SiN membranes are also flexible and can bow while spinning, which adds additional variability and irreproducibility. If the substrate is not perfectly centered in the chuck it will also cause uneven resist. It is important to inspect the substrates after spinning to make sure the resist is evenly coated. Sometimes the first layer of resist can be re-spun if it does not come out correctly. Certain failure mechanisms in sample fabrication could be caused by uneven or irreproducible resist. These include poor liftoff, due to there not being enough undercut, or individual elements being too thin or too wide due to wrong thickness of resist.

After the layers of resist are coated and baked and the aquaSAVE is spun on, the sample are ready to be loaded into the vacuum chamber for lithography. The electron beam is carefully aligned, stigmated, and focused on the surface of the samples. I use a 10 kV accelerating potential, a 30 µm aperture for the beam, and a 10 mm working distance between the gun and sample. The accelerating potential can affect the resolution and size

of elements written. Typically, higher energy means smaller features, but ~100 nm wide features do not require a higher potential and I found 10 kV to be sufficient. The aperture size sets the beam current. I always used a 30 μm aperture, but a smaller aperture might help reduce disorder in future work. The electron beam moves extremely fast during patterning. I observed some effects of elements bending or missing in finished samples that may be reduced if the beam moves slower. A smaller aperture will reduce the beam current and if the electron dose is kept the same, the beam will have to move slower to deliver the same dose with lower current.

After focusing and aligning the beam, the beam current is measured and the electron dose must be specified. The doses for square and kagome geometries were carefully calibrated through dose arrays discussed before as well as trial and error. The doses are inherently related to and dependent on the exact geometry and design of the specific crystal being written. Connected kagome structures have connected lines in the digital pattern and require certain doses based on the length of the lines to get the widths right. In the square geometry, I want the elements disconnected and so the lines in the digital pattern are disconnected as well. In the kagome geometry, I fabricated connected samples of varying element length. There are proximity effects in lithography so that shorter elements need lower doses in order to have similar widths to the longer ones with higher doses. I tried to correct for this as much as possible, but sometimes shorter elements I fabricated still tended to be wider and this was an important consideration when modeling the kagome ordering behaviors discussed in Chapter 3. Typical doses for 500 nm length connected kagome elements that result in ~100 nm wide magnets are around 600 pC cm^{-1}. For the shortest 300 nm length elements, the dose was closer to 350 or 400 pC cm^{-1} and this resulted in

elements around 130 nm wide. For the square geometry elements and topological defect crystals in Chapter 4 (which are derivations of the square geometry) I use a dose of 550 pC cm^{-1}. For the square lattices I try to get the elements as close as possible without touching to increase magnetostatic coupling between neighboring elements. This is done through a combination of the electron dose and the actual length of the line in the patterning software. I typically used a lattice spacing of 500 nm with a line length of 340 nm in the software, which resulted in magnets with an actual length around 400-420 nm long. The square geometry elements were typically 110-120 nm wide as well.

After pattering, the resist needs to be developed. First, the samples are placed in de-ionized water to remove the aquaSAVE layer. The samples are then submerged in methyl isobutyl ketone (MIBK) developer mixed with isopropanol (IPA) in a ratio of 1:3 MIBK:IPA for 1 minute. The samples are then transferred to pure IPA for 1 minute to stop the development process. Finally, they are removed and blown dry with pure nitrogen gas.

2.2.2 Electron beam evaporation

Following development, the next step is to electron beam evaporate an Al etch mask over the patterned resist. Some of the Al will land on the resist, but wherever the electron beam was traced, the Al will land directly on the $FePd_3$ as shown in Figure 2.11(e). Electron beam evaporation is a common and versatile method to deposit a wide range of materials. It is quite similar to thermal evaporation and both techniques work through the same principle. A material is heated inside a high vacuum chamber until the material's vapor pressure becomes comparable to the chamber pressure. The vapor pressure increases with temperature, so as the material gets hotter, its vapor pressure rises until it equals the

ambient pressure and the atoms in the material start to evaporate. Typically, materials with higher melting points require higher temperatures in order to start evaporating.

Both thermal and electron beam evaporation use this process, they differ in how the material is actually heated. In thermal evaporation, a small pellet of desired material is placed in a conducting boat, typically tungsten, and current is passed through the boat which heats due to resistive, or joule, heating. Electron beam evaporation uses a high energy beam of electrons focused at a small crucible containing the desired material. The electrons heat the material until it melts and starts to evaporate. Electron beam evaporation can get materials hotter because the energy and high temperature is in a more focused region. This allows for the evaporation of a lot of oxides and other high melting point materials and makes electron beam evaporation more versatile than thermal evaporation, though both techniques work equally as well for some materials. Generally, for both techniques, the source material has to be at the bottom of the vacuum chamber in order for the boat or crucible to hold it and the sample is placed upside down at the top of the chamber, shown schematically for electron beam evaporation in Figure 2.14.

I always used electron beam evaporation, typically in a Denton tool at UMD, to evaporate my films. As was discussed at the beginning of Section 2.2, evaporation is essential for the liftoff process I use after lithography, but it is also a quick, easy, and reliable way to deposit the underlayers for my $FePd_3$ films as well. The Denton has multiple material pockets so that I can deposit both underlayers consecutively without having to break vacuum. No matter what material I am depositing, I clip my SiN substrates into a custom deposition holder (see Figure 2.15) so that they can be mounted upside down inside

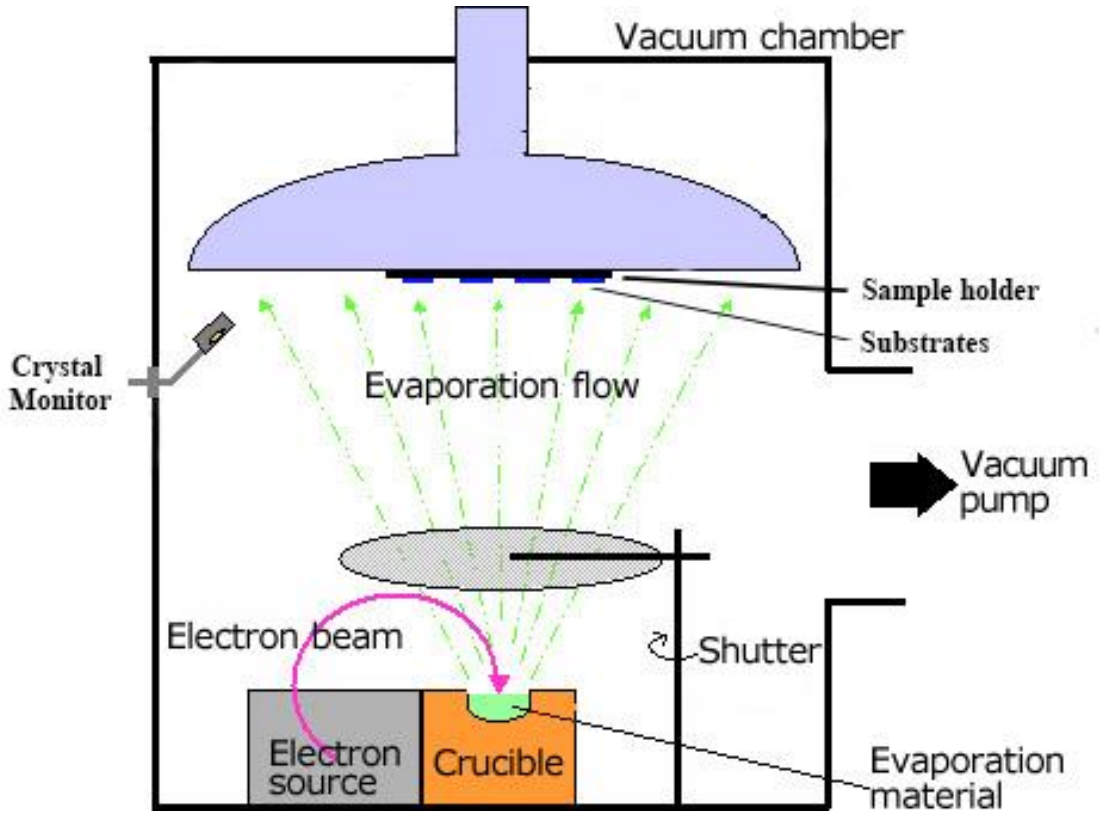

Figure 2.14 **Electron beam evaporation**. Electrons are directed into a crucible with evaporation material inside a vacuum chamber. The materials heats up, melts, and starts evaporating. The atoms land on the substrates mounted upside down at the top of the chamber. A quartz crystal monitor measures the material deposition rate and thickness. Figure adapted from Ref. [178].

the Denton chamber. I carefully calibrated the crystal monitor just like in the sputter chamber in Section 2.1.2 to evaporate precisely known thicknesses of the different materials.

To deposit the Al etch mask for sample patterning (Figure 2.11(e)) I first load the developed samples into the vacuum chamber and pump it down. Once a base pressure of around 3×10^{-6} Torr is reached, the electron beam is turned on and aimed at the Al crucible. The beam power is increased until the metal begins to melt. It is important to turn on the beam modulation so that the metal melts evenly and the crucible will not be damaged. Once the deposition rate reaches about 0.4 Å s^{-1}, I open the shutter covering the samples

Figure 2.15 **Sample holder**. Optical photograph of SiN substrate samples clipped into deposition and milling holder. The frame is made from Al and the clips are Be/Cu. The holder is colored due to various materials depositions. Scale bar is 0.5 inches.

and begin the deposition. 16 nm of Al is then deposited at a rate of 0.4-0.8 Å s^{-1}. Once 16 nm is reached, the shutter is moved back in front of the samples and the electron beam power is ramped down and turned off. The Al crucible is still extremely hot and must be allowed to cool for 5-10 minutes before the chamber is vented and the samples are removed.

2.2.3 Liftoff

Liftoff is a common process used in the fabrication of many types of structures. It works through a solvent removing the resist left on a substrate after lithography and metal deposition, causing the metal deposited on top of the resist to “lift off” as shown in Figure 2.11(e)-(f). It helps to have an undercut in the resist so that the solvent can readily access

the resist to remove it. Liftoff is generally easy, though using it with SiN membranes adds some challenges to the technique. Typically, samples are agitated with ultrasonic waves or "sonicated" during liftoff to assist the solvent with removing the resist and metal. The problem with doing this to the membranes is that the extremely thin membrane will immediately break if it is subjected to sonication. Instead, I developed other techniques to aid the liftoff process, though it is still one of the most difficult steps in my fabrication, especially for the kagome crystals.

The kagome crystals are much harder to liftoff because of their connected structure. In the square geometry samples, the elements are disconnected, so the Al film on top of the resist is nominally continuous and can come off in large flakes. In a kagome crystal, every connected hexagon of desired nanomagnets has its own hexagonal piece of material that needs to be lifted off in the middle. We have named these pieces "cookies" since they look like little cookies on the sample. It can be extremely difficult to get these cookies to liftoff. Sometimes large areas of a membrane will not liftoff at all due to problems in spinning or developing the resist, but even if the sample has generally good liftoff, there were almost always a few cookies left on the sample as defects. An image of some cookies is shown in Figure 2.19 in Section 2.2.5.

To liftoff the Al on top of the PMMA I first submerge the samples in an acetone bath and soak them for 1 hour. The next step is a trick I have developed and perfected over time for the membranes. I take a spray bottle filled with acetone in one hand and grab a sample with tweezers in the other hand. Keeping the sample submerged at first, I begin to spray the top of the acetone bath and I lift the sample out of the bath and into the continuous spray. The continuous spray breaks the surface tension of the small droplet of acetone that

normally forms on top of the membrane if you just lift it out of the bath as is. I have found that this trick greatly helps reduce the number of cookies left on the kagome samples and generally results in better liftoff than just lifting the samples out, allowing the droplet to form, and then spraying them. Spraying the samples with a strong stream of acetone is also very important and works something like sonication during liftoff. The spray helps to forcibly remove the resist and metal from the membrane, but without breaking the SiN membrane. It is important to spray from all angles to get the best results.

After spraying, the samples are placed back in the acetone bath for 20 minutes for the square geometry sample and 1 hour for the kagome geometry samples. All samples are then sprayed again, using the same technique to break the surface tension as they are removed from the bath. Finally, the samples are submerged in IPA for 2-3 minutes to help remove any residue from the acetone. They are removed from the IPA and blown dry with nitrogen gas.

2.2.4 Ion milling

After liftoff, we are left with patterned Al elements on top of the $FePd_3$ film as shown in Figure 2.11(f). The next step is to etch the entire sample using Ar^+ ion milling. Ion milling is very similar to sputtering in that an Ar plasma is created inside a vacuum chamber and the Ar^+ ions are accelerated to knock material off of a target. The difference here is that in ion milling the Ar is being accelerated towards the sample and no material is being deposited anywhere, only removed. The Ar atoms bombard the whole surface of the sample and etch the top layer which is either the patterned Al or the $FePd_3$ film. We want to etch away all of the $FePd_3$ everywhere it is not protected by the Al and leave the

$FePd_3$ that is under the Al in order to transfer the pattern of the Al into the $FePd_3$ film, as shown in Figure 2.11(g)-(h).

The Ar atoms etch the Al and $FePd_3$ at different rates. I had to carefully calibrate the etch rates simultaneously with the $FePd_3$ sputter deposition thickness calibration discussed in Section 2.1.2. Using an ion mill in CNAM at UMD, I would etch a film of known thickness with some drops of resist baked on the substrate for 1 minute at 200 W power. I then remove the resist and material on top of it and use a profilometer to measure the step height. Then I compare this height to the measured, non-etched thickness of the original film and use the difference in the film thicknesses to calculate an etch rate. The step height of the etched films tends to be difficult to measure, but by averaging over many samples and many points on each sample, I was able to obtain a consistent and accurate rate. I found the Al etches at a rate of 15.3 nm min^{-1} and the $FePd_3$ at a rate of 32.5 nm min^{-1}. The difference in the rates is striking, but comes from the Pd being a high higher Z (atomic number) material than Al, which actually etch faster than lower Z materials.

At 32.5 nm min^{-1} it should take ~42 seconds to etch through the 23 nm $FePd_3$. I typically etched for slightly longer, 45-50 seconds, to make sure all of the material would be removed. I don't want to etch much longer than 50 seconds because then the SiN membrane will start to be etched which could weaken or break it. Luckily I never saw this problem in my samples. I use the 45 second etch time to calculate how much Al to deposit during electron beam evaporation. Based on the calibrated rate, 11.5 nm of Al will be etched in 45 seconds, so I deposit a little more (16 nm) to make sure none of the $FePd_3$ I want to keep is etched away. A few extra nm of Al on top of the $FePd_3$ elements will also help protect them from oxidation and it does not affect their magnetic properties.

After ion milling, the desired pattern has been etched into the $FePd_3$ film, shown schematically in Figure 2.11(h). There is a final step to electron beam evaporate a 5 nm Ti layer over the entire sample. This layer acts as an anti-charging layer for TEM studies. Since the SiN is insulating, electrons can build up on the surface, creating an electric field that will deflect or move the electron beam. This can cause images to shift around and make it impossible to take clear micrographs. The Ti layer also helps protect the $FePd_3$ elements from oxidation.

Images of some finished samples are shown in Figures 2.16, 2.17, and 2.18. The many different samples and geometries I fabricated will be discussed in later chapters, but these figures provide representative examples of the two main geometries (square and kagome). Figure 2.16 shows a full square geometry sample of 420 nm x 110 nm elements with a 500 nm lattice spacing. Figure 2.17 shows a large kagome crystal along with some smaller kagome crystals. There are some one and two hexagon crystals in the upper right corner and some medium sized crystals with different edge types in the bottom right. The different types of edges and their relative effects will be discussed in depth in Chapter 3. Figure 2.18 shows higher magnification images of the square and kagome geometries showing more detail of the elements.

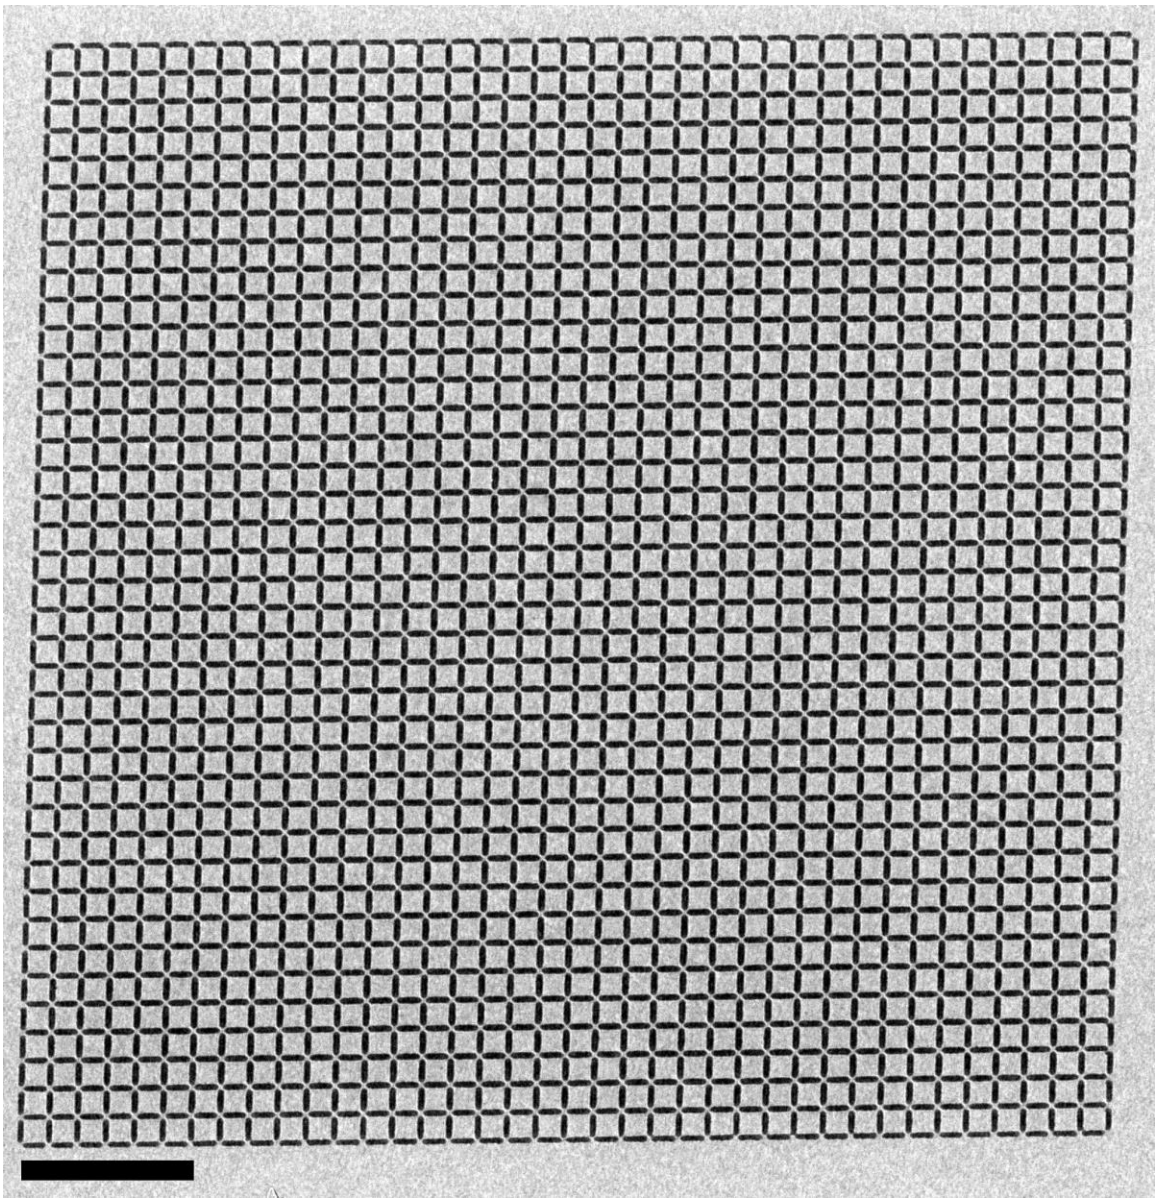

Figure 2.16 **Large square crystal**. In-focus TEM image of a square geometry sample. Scale bar is 3 μm.

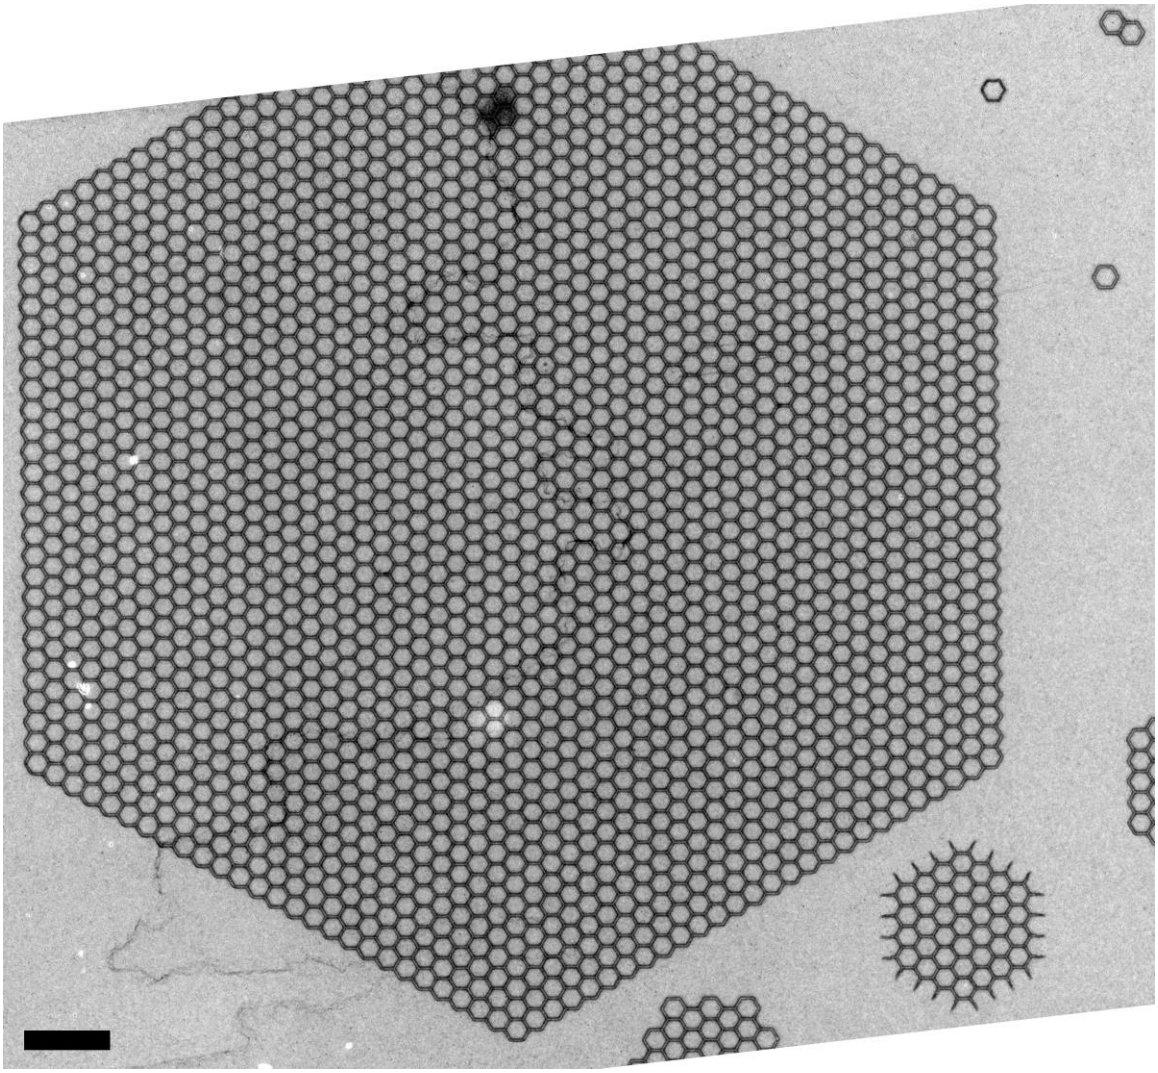

Figure 2.17 **Large kagome crystal**. Lorentz TEM image of a large kagome crystal along with a few smaller crystals. The crystals in the top right are one and two-hexagon crystals and there are three medium sized crystals in the bottom right. Scale bar is 3 μm.

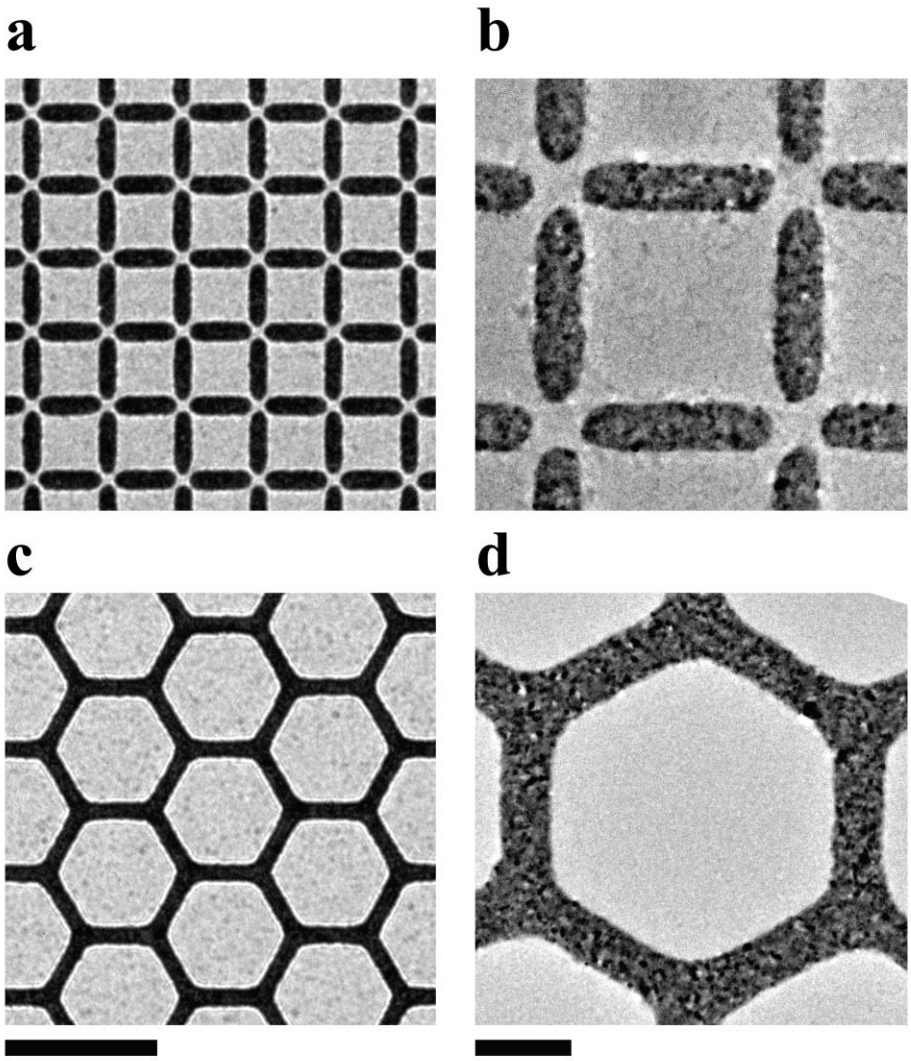

Figure 2.18 **Higher magnification square and kagome images**. **a**. In-focus TEM image of a square geometry sample at moderate magnification. **b**. High-resolution, high-magnification TEM image of square geometry elements. The grain structure of the $FePd_3$ is visible. **c**. In-focus TEM image of a kagome geometry sample at moderate magnification. **d**. High-resolution, high-magnification TEM image of connected kagome elements. Scale bar for (**a**) and (**c**) is 1 μm, scale bar for (**b**) and (**d**) is 200 nm.

2.2.5 Common problems

The most common problems I encountered during fabrication are membranes breaking and poor liftoff. The membranes are handled frequently during the fabrication process, from spin-coating the resist to development to loading and unloading the

deposition holder, liftoff, and more. If the tweezer does not have a solid grip on the sample or if it has too strong a grip, the membrane can drop or fly out and break. The membrane can also be launched out of the spin coater chuck when it is spinning at high speeds. Careful focus and practice can help prevent most of the membrane breaking problems, though some loss is probably inevitable. The rest of this section discusses other common issues with fabrication, plausible causes for these issues, and ways to help prevent these issues in future work.

2.2.5.1 Poor liftoff

Poor liftoff is a difficult issue to overcome. I developed some tricks that seemed to help, but there are always some areas of a sample that do not lift off well. Many times crystals near the edges of the membrane do not liftoff as well, probably because of variations in the resist near the edges. There are regularly cookies in even the best kagome crystals. Figure 2.19(a) shows a kagome crystal with a large number of cookies on it. This crystal has too many cookies to be usable, though good data and insight can still be gained from crystals that have just a few cookies. Figure 2.20 shows a large region of a sample where some kagome crystals lifted off and others did not. There is some dependence on size and dose, as smaller element length kagome crystals have a harder time lifting off because the cookies are smaller, but there are still some crystals of the same size and dose that have different degrees of success with lifting off.

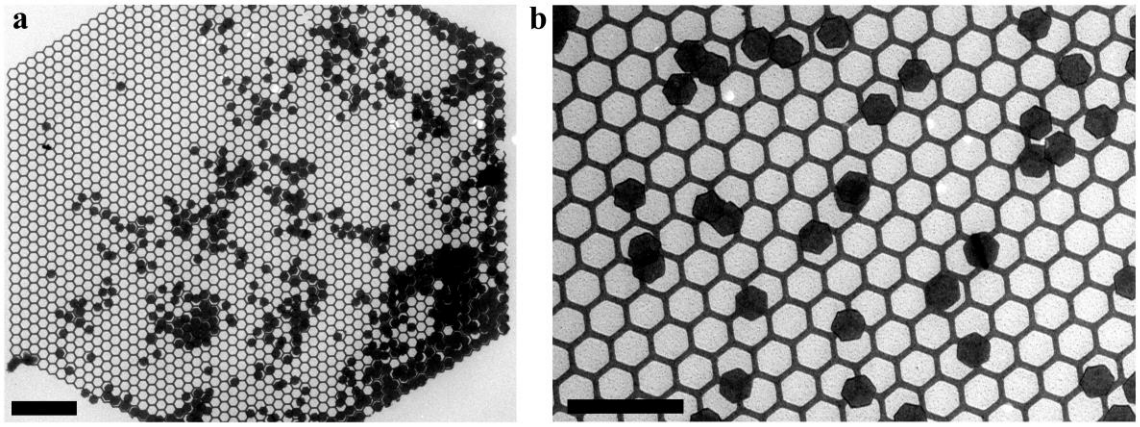

Figure 2.19 **Kagome geometry with cookies**. **a**. Low magnification, in-focus TEM image of a kagome crystal with lots of cookies present, representing poor liftoff. Scale bar is 5 μm. **b**. Higher magnification TEM image of the cookies showing more detail, in particular their hexagonal shape. Scale bar is 2 μm.

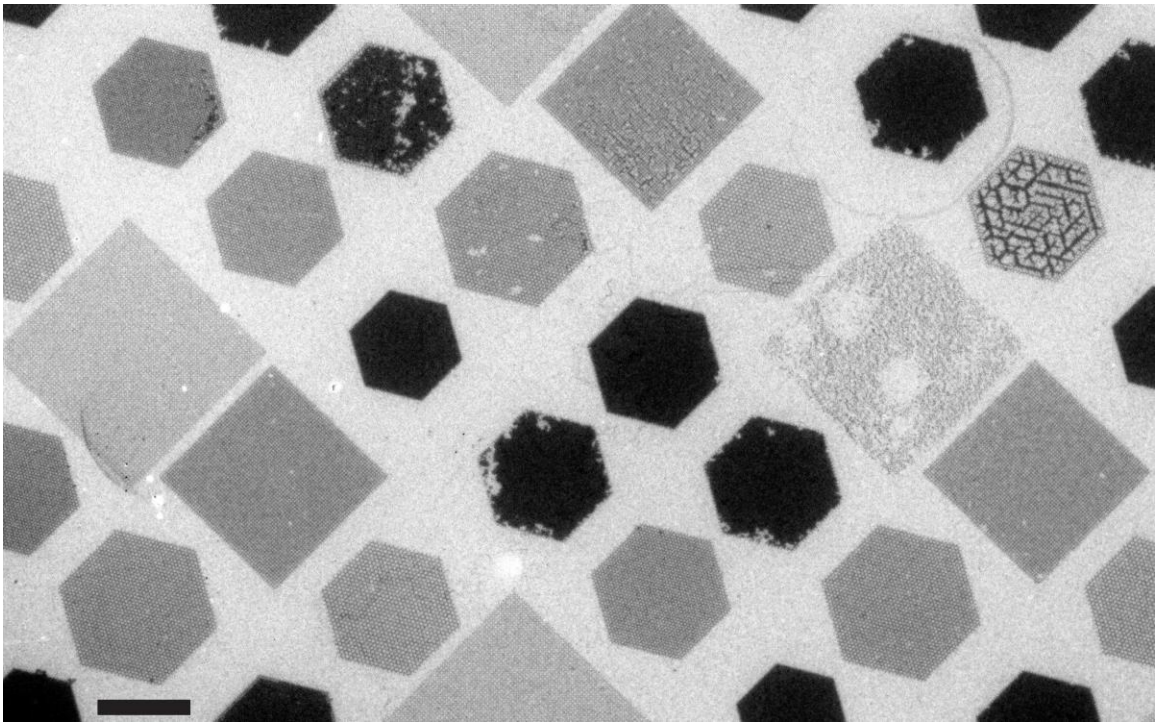

Figure 2.20 **Sample with poor liftoff**. TEM image of a large region of a sample with both square and kagome geometry crystals. Some kagome crystals show good liftoff, some have poor liftoff, and others have no liftoff at all. There are also many elements missing from one of the square geometry crystals. Scale bar is 10 μm.

2.2.5.2 Spin coating resist and electron dose

There are a variety of problems I encountered related to the thickness and consistency of spin coating PMMA and local electron dose effects, though these tend to affect individual crystals or smaller regions of a sample more than poor liftoff or membrane breakage. One prevalent issue is that individual elements will be missing from a larger array. This tends to be more common in the square geometry samples because the elements are disconnected, but it is found in the kagome geometry as well. Examples are shown in Figure 2.21. Missing elements may be due to local defects in the resist due to artifacts from spin coating or from spraying the membranes too hard with acetone during liftoff. This could be prevented by spraying less and more gently, though there is a tradeoff because some spraying is generally necessary to aid in liftoff.

Elements in the square geometry can also be bent or deformed as is shown in Figure 2.22. This is likely due to one or more layers of resist not spinning on at a uniform thickness, altering the optimal electron doses for different regions of the membrane. The bent elements may also be caused by the electron beam charging the surface of the sample if the aquaSAVE layer did not spin on evenly. Sometimes dirt or chemical residue can be left on a crystal as well, masking or covering some of the elements. Finally, the elements can come out the wrong width, either too thin or too wide. This can be caused by improper dose calibration during lithography, an unknown change in the beam current during lithography, or the resist spinning on at a different thickness than during the dose testing. Sometimes different regions of a sample can have different widths which is likely caused by variations in the thickness of the resist across the sample. An example of both under- and overdosed elements is shown in Figure 2.23.

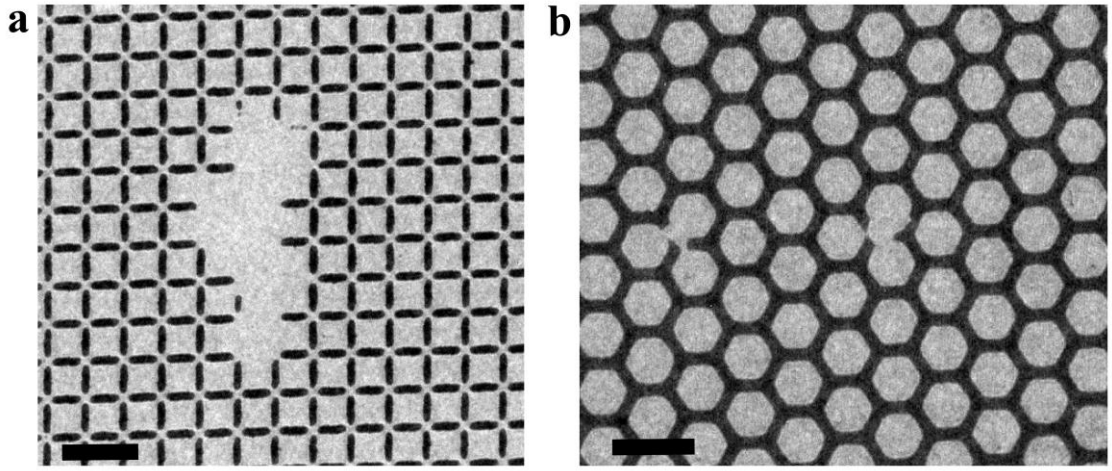

Figure 2.21 **Missing elements**. **a**. In-focus TEM image of a square geometry sample with a region of missing elements. **b**. In-focus TEM image of a kagome sample with only a few missing or broken elements. Such defects do not substantially affect the long-range ordering in these samples. Scale bars are 1 μm.

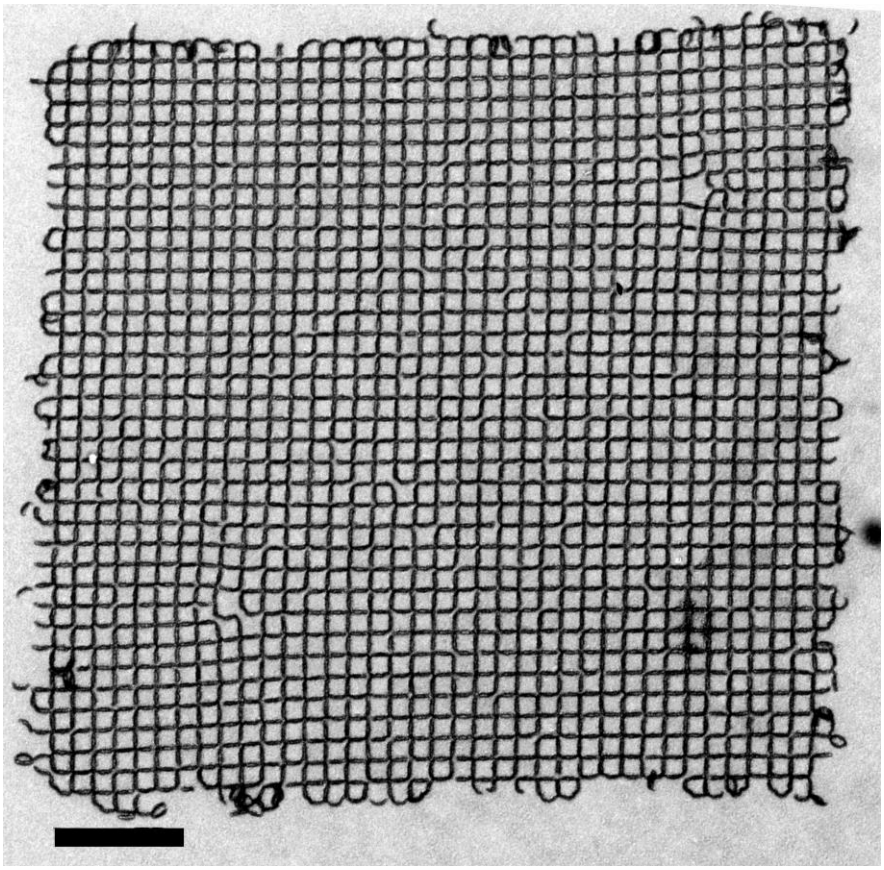

Figure 2.22 **Bent elements**. In-focus TEM image of a topological defect square geometry sample. Many elements are touching or curved. Scale bar is 3 μm.

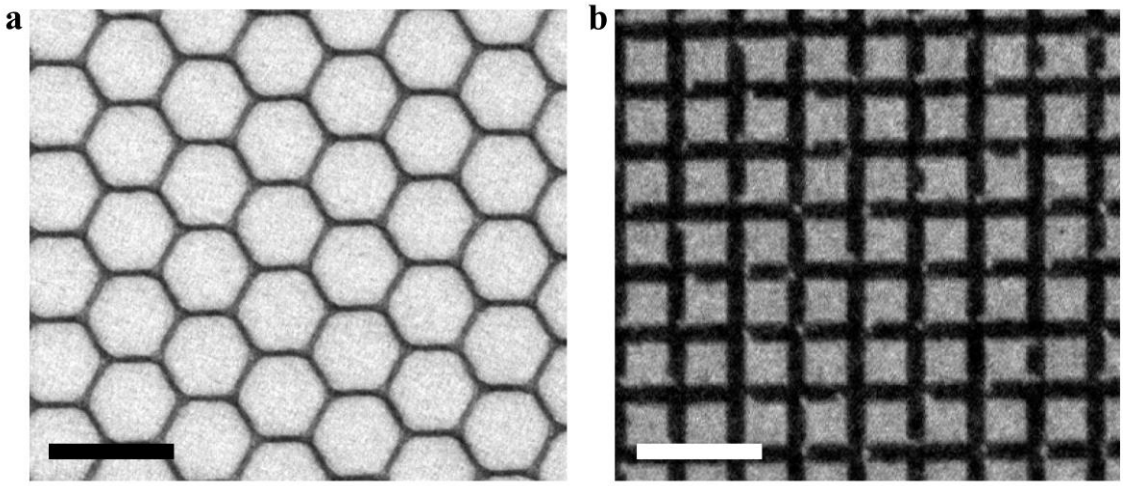

Figure 2.23 **Wrong doses**. **a**. In-focus TEM image of a kagome sample with too low a dose. The elements are extremely thin and too small. **b**. TEM image of a square geometry sample with too high a dose. The elements are supposed to be disconnected, but came out connected and too wide because the dose was too high. Scale bars are 1 μm.

2.2.5.3 Membrane quality

A major issue I ran into at one point in my work related to the quality of the SiN membranes themselves. We purchase the membranes commercially and discovered that a large batch we had received has dark spots in them (see Figure 2.24(a)). The membranes should not add any extra contrast to the TEM images and so the dark spots posed a significant problem. We believe that the spots are Si-rich regions within the SiN. The manufacturer can fabricate SiN that is close to the Si_3N_4 stoichiometric ratio which has no spots, but is more fragile, or they can make the films Si-rich, which gives them more strength, but also these dark spots. The main problem with the dark Si spots is that in the Lorentz TEM mode (see Section 2.3.2) that we use to image the precise magnetic configurations of the samples, the dark spots become very bright spots (see Figure 2.24(b)). The magnetization of each element is determined by very slight asymmetry in the contrast across that element. We use image processing software (see Section 2.3.3) to extract the

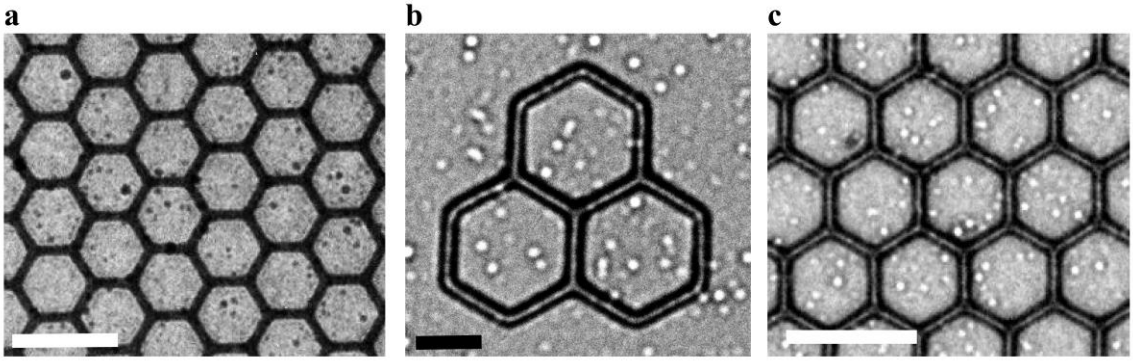

Figure 2.24 **Spots on SiN membranes**. **a**. In-focus TEM image of a section of a kagome crystal showing SiN membrane spots. The spots appear dark when in focus. Scale bar is 1 μm. **b**. Lorentz contrast TEM image of a 3-hexagon kagome crystal showing the SiN membrane spots. The spots appear bright in the Lorentz imaging mode. Scale bar is 500 nm. **c**. Lorentz contrast TEM image of a section of a large kagome crystal also showing the bright spots. The spots skew the contrast asymmetry across each element that normally determines the magnetization direction. The bright spots cause the image processing software to incorrectly assign spin directions, making the data from these samples unusable. See Section 2.3.3 for a further discussion. Scale bar is 1 μm.

magnetization direction of the individual magnets. The bright spots cause our image processing software to read many of the magnetization directions incorrectly because the bright spots skew the contrast asymmetry. This results in analysis that is wrong, not representative of the actual magnetic configuration, and completely unusable. After this problem was discovered, new samples had to be fabricated on new membranes that did not have the spots.

2.3 Transmission electron microscopy

2.3.1 Fundamentals and advantages

The transmission electron microscope is the main experimental tool I use to image my ASI structures. Transmission electron microscopy (TEM) is an extremely versatile and powerful technique for analyzing all kinds of different materials. The microscopes work

by accelerating electrons at extremely high energies down a vacuum column, transmitting them through a specimen, and collecting them on the other side. The specimen blocks or diffracts the electrons attempting to pass through and the variation in the number of the electrons hitting different parts of a florescent screen or CCD camera creates contrast differences that give information about the specimen being studied.

There are many different TEM imaging modes that can provide different types of information about a specimen. High-resolution or atomic-resolution imaging is very common and a particular advantage of TEM. The de Broglie wavelength of an electron in a typical 200 kV microscope is on the order of 2.5 pm [175]. This allows for some of the best atomic-resolution images possible in a wide variety of materials. Lorentz TEM is the technique I use most and will be discussed at length in the next section, but in general it can provide magnetic information about a specimen. There are also multiple techniques for elemental analysis built into many microscopes. EDS was discussed in Section 2.1.3 and there are often EDS detectors inside the microscope column that can provide qualitative information on what elements are present in a specimen. Electron Energy Loss Spectroscopy (EELS) can give not only elemental information, but also specific atomic bonding information and can tell the difference between different molecules and molecular bonding in a substance. EELS can be combined with atomic resolution imaging for incredibly precise investigations of crystal structure and atomic arrangements. This also allows for novel investigations of interfaces and defects because the exact atomic positions and species are known. Electron diffraction is a powerful technique for revealing the crystal structure of many different kinds of materials. In some cases it is even better than x-ray diffraction because the electron wavelength is smaller than the wavelength of x-rays and

so the electrons are sensitive to smaller features and more subtle differences between the structures of some materials or different phases of the same material. Normal TEM imaging is done in what is called “bright field” mode, where higher Z elements appear darker in the CCD images. “Dark field” imaging is also possible, which uses diffracted electrons to image the sample and can reveal dislocations and grain boundaries very easily, among other information. An important factor in all of these imaging modes is that the specimens have to be thin in order for electrons to transmit through them, typically around 100 nm or less. Thicker samples can be milled down using ion milling or Focused Ion Beam (FIB) in order to prepare them for TEM investigation.

A schematic of a transmission electron microscope is shown in Figure 2.25. Electrons are generated at the top of a vacuum column and accelerated with an electrostatic potential typically of around 200 kV, though some tools are capable of imaging at as low as 50 kV or as high as 1 MV. The beam of electrons then passes through a series of apertures and lenses as well as through the sample itself before being projected onto a florescent screen or a CCD camera. The CCD camera exposure can be adjusted to integrate electrons reaching it over a period of time and provide better contrast images. The lenses are complex electromagnets, or solenoids, though their principles of operation are very similar to the glass lenses in an optical microscope. A ray diagram for the electrons is shown in Figure 2.26. Electrons bend in a magnetic field due to the Lorentz force, so the magnetic fields that the solenoids generate can spread or focus the beam of electrons as well as shift its position. Before imaging, the whole instrument is aligned in order to capture high quality images. The alignment procedures generally differ for the different imaging modes.

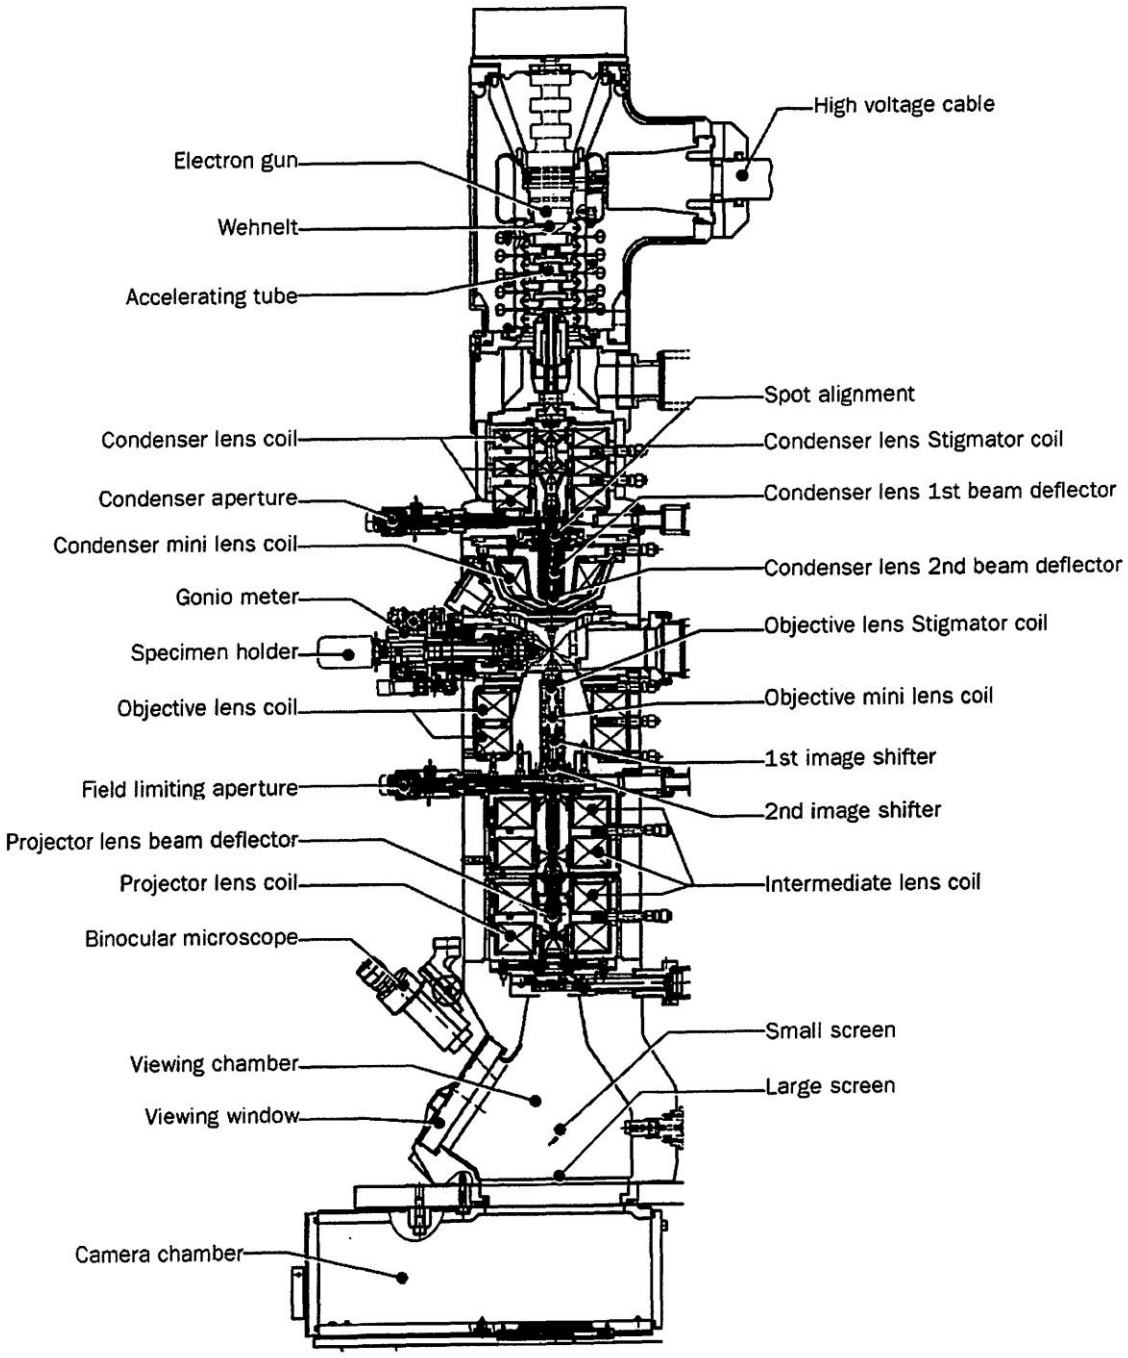

Figure 2.25 **Schematic of a TEM**. Cross sectional view of a TEM column showing the locations and relative sizes of the parts of the microscope. Figure adapted from Ref. [179].

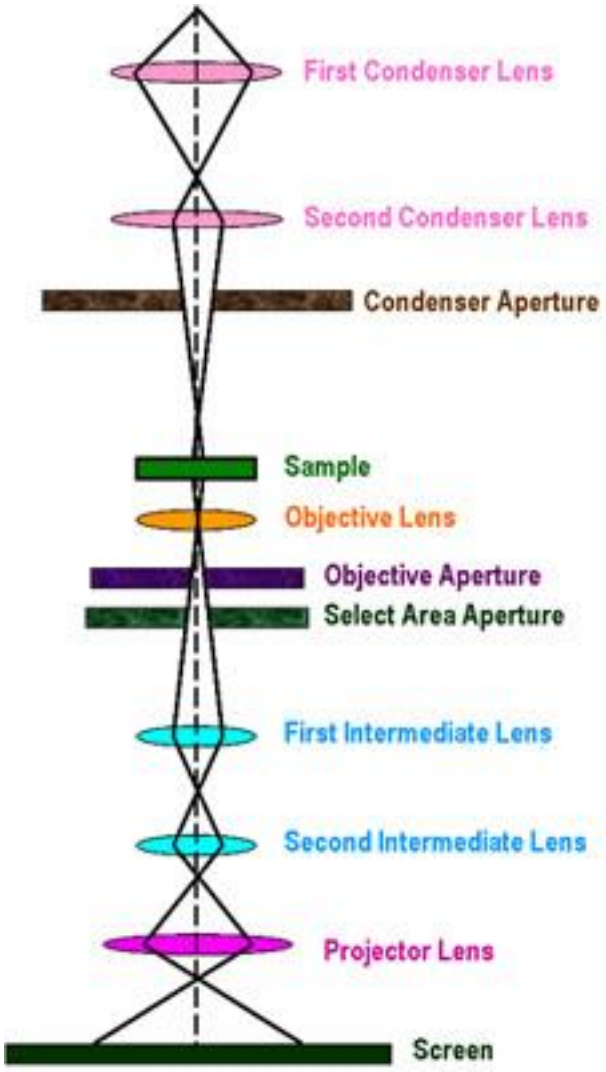

Figure 2.26 **TEM ray diagram**. Schematic of electron trajectories in a TEM and how the various lenses and apertures affect the electron beam. Figure adapted from Ref. [180].

The capabilities and resolution limits of a microscope are determined by a number of factors. Recently, some of the biggest improvements to imaging have come through aberration correctors. There are two types of aberrations, chromatic and spherical. Chromatic aberration is caused by the electrons in a beam not having the exact same energy, but instead a spread of energies, usually Gaussian, centered around the nominal accelerating potential. The different energy electrons will respond differently to a lens,

preventing the lens from being able to perfectly focus the whole beam and limiting the minimum resolution. Spherical aberration arise because electrons passing through the middle of the lens are focused differently from electrons passing through the edge of the lens. In an optical microscope it is possible to correct for this by using a divergent lens, but for a long time the solenoid TEM lenses could only be made convergent and thus could not correct for spherical aberration. New spherical aberration correctors effectively act as divergent electromagnetic lenses and have enabled new regimes of TEM studies. New gun sources can create more uniform energy distributions, reducing the chromatic aberration as well. A final limiting factor for TEM studies is beam damage. The high energy beam of electrons interacts with every specimen and can cause damage, especially to biological or very small samples. Reducing the accelerating potential of the TEM can help alleviate this problem, though it can never be completely removed. Many of the newest microscopes combine low kV imaging, to prevent damage to sensitive samples, with aberration correction to achieve the best possible images.

The microscope I used most for my analysis was a JEOL JEM 2100 LaB6 tool operating at 200 kV. It worked extremely well for taking all the images needed for this dissertation. The "LaB6" refers to how the electrons are generated in this microscope; a lanthanum hexaboride (LaB_6) filament generates the beam of electrons. Most of the imaging I did was in low-magnification mode (< 2000x), which is ideal for imaging arrays with thousands of ASI elements, both in-focus and in Lorentz mode which will be discussed in the next section (see Figures 2.16 and 2.35 for examples). The LaB6 microscope has a large gap inside the objective lens which allows for more versatile imaging, better low-magnification imaging, and a variety of specialized sample holders to

fit inside. All members of Cumings' group use *in-situ* TEM to study various nanoscale phenomena, apply stimuli inside the microscope, and observe what is happening on the nanoscale in real time. I most often used a specialized heating holder, which could heat samples up to 500 °C or more inside the microscope. I used this to measure the T_C of $FePd_3$ films as well as do some of the initial heating of my patterned thermally active ASI samples. We also have some custom electronics that let us take over control of the objective lens, the most powerful lens in the microscope and the one that is directly around the specimen. Low-magnification imaging does not require use of the objective lens and so it is normally "off" in this mode. We are able to switch over control of the lens to an external power supply that is controlled by a LabVIEW program, allowing us to apply arbitrary magnetic fields *in-situ* or remove any remnant field of the objective lens that could bias a sensitive sample. The LabVIEW software can also automate some other tasks in the microscope such as taking images, although I did not use this often in my studies. Finally, I did do some high-resolution TEM imaging to both investigate the grain structure of my $FePd_3$ films (see Figure 2.6(b)) as well as inspecting and measuring the size of finished, patterned ASI elements (see Figure 2.18(b) and (d)). Overall, the transmission electron microscope is an extremely powerful tool that was crucial for all parts of my dissertation work.

2.3.2 Lorentz TEM

Lorentz TEM is a technique that exploits the Lorentz force for gaining magnetic information about a sample. A schematic of the technique is shown in Figure 2.27. As electrons are transmitted through a material with in-plane magnetization, the electrons are

deflected, due to the Lorentz force from the magnetic field inside the material, in different directions based on the direction of the in-plane magnetization. This can create added bright or dark contrast in the TEM images that provides information about the magnetic structure or properties of the specimen. This contrast can be seen as domain walls and ripple contrast in continuous thin films found in Figures 1.9, 1.10, 2.7, and 2.8. In the ASI nanomagnets, imaging in Lorentz mode creates a fairly subtle asymmetry in the contrast across each element. This can be seen in Figure 2.28(b) and (d). There is a darker black line on one side of the element then on the other side. The line profile in Figure 2.28(e) shows the asymmetry more clearly and shows that it differs for the two different magnetization directions.

Yi Qi, a graduate student in Cumings' group who worked on the ASI project before Daunheimer, ran some simulations of the Lorentz TEM technique. He used a contrast transfer function to confirm and interpret the observed contrast differences in Lorentz TEM images were due to a single magnetic domain oriented in different directions [181]. His simulation results are shown in Figure 2.29. A full explanation of the fundamental origins and spacial variation of Lorentz contrast in arbitrary geometries is somewhat involved and beyond the scope of this dissertation. Derivations can be found, for example, in Refs. [181] or [182]. It will suffice to note that the observed contrast is phase contrast from interference of electron plane waves with phase shifts due to the magnetic field inside the specimen. The phase shift in the plane wave of electrons due to the in-plane magnetic field is given by

$$\varphi(x) = -\frac{2\pi e}{h} Mtx \tag{2.3}$$

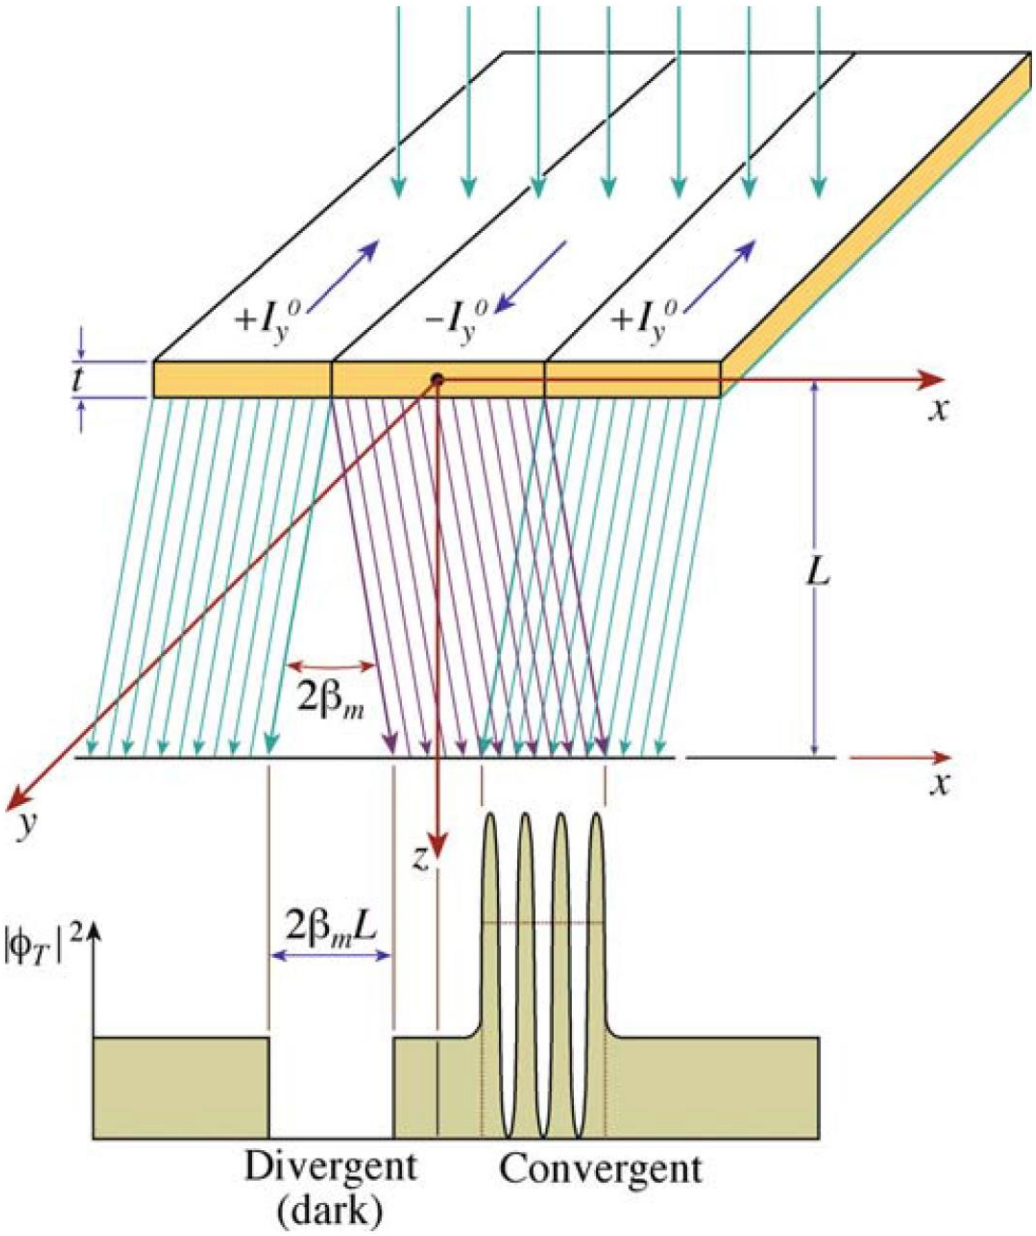

Figure 2.27 **Lorentz TEM**. As electrons transmit through a material with in-plane magnetization, the electrons are deflected, due to the Lorentz force, in one direction or another based on the direction of the magnetization. This can lead to added bright or dark contrast due to interference fringes in the TEM image. Figure adapted from Ref. [175].

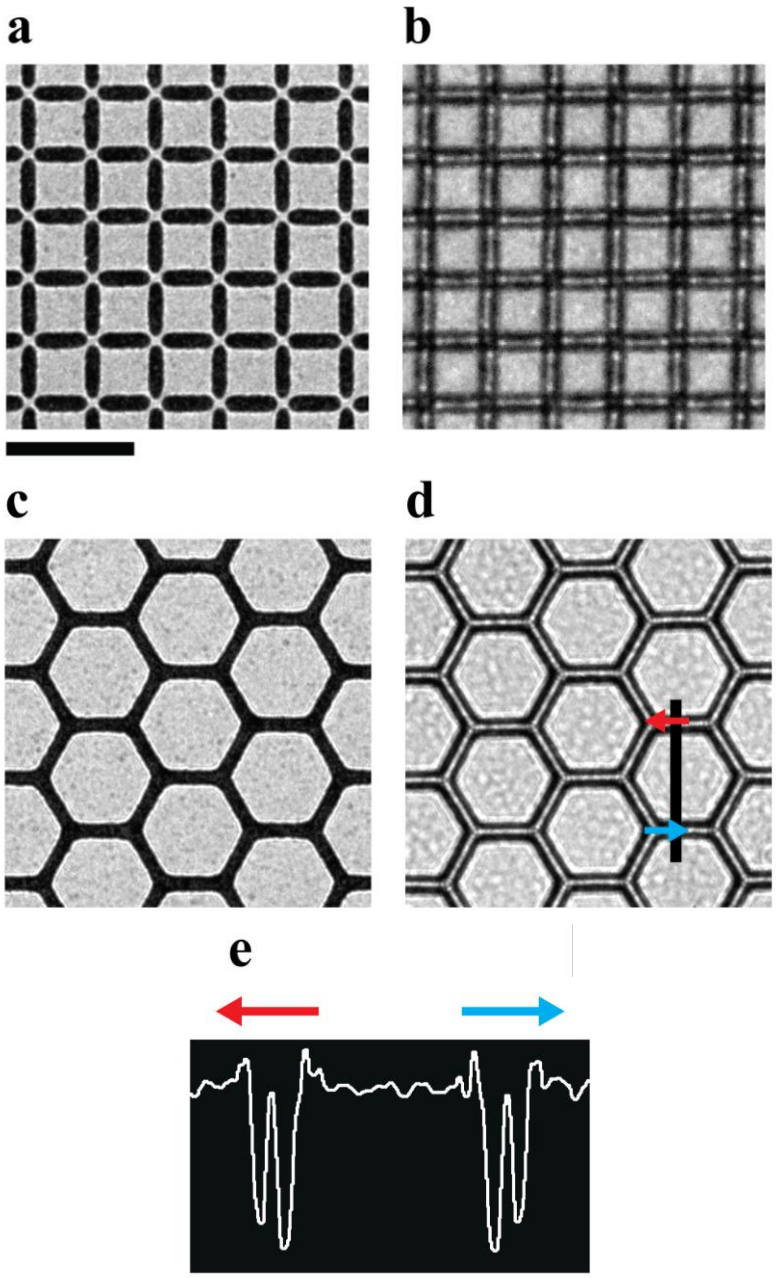

Figure 2.28 **In-focus vs. Lorentz contrast TEM images**. **a**. In-focus TEM image of a square geometry sample. **b**. Lorentz contrast TEM image of the same square geometry sample, showing ground-state order. **c**. In-focus TEM image of a kagome geometry sample. **d**. Lorentz contrast TEM image of the same kagome geometry sample, showing only ice-rule ordering. **e**. Intensity profile along the black line in (**d**) with corresponding arrows showing how the asymmetry in contrast across each element is used to determine the direction of the magnetic moment. Scale bar is 1 μm.

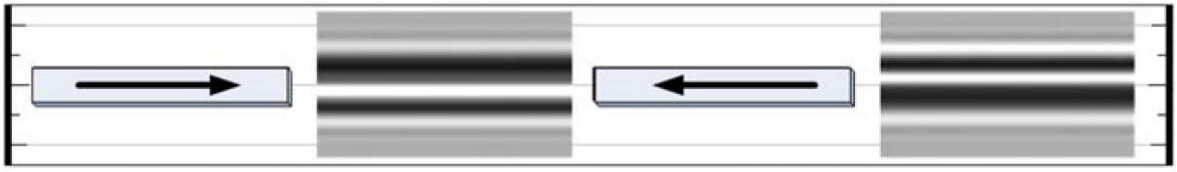

Figure 2.29 **Lorentz contrast simulation**. Contrast transfer function simulation confirming that the observed contrast asymmetry in Lorentz TEM images is due to magnetization of individual elements in different directions. Figure adapted from Ref. [93].

where *e* is the electron charge, *h* is Plank's constant, *M* is the magnetization, *t* is the thickness of the film, and *x* is the spacial coordinate in the object plane perpendicular to the direction of **M**. The contrast intensity distribution is then found to be [181]

$$I(x',y') = \left|\iint\left(\iint \exp[i\varphi(x)]\exp\left[-2\pi i\left(k_x x + k_y y\right)\right]dxdy\right) * \exp[\pi i\Delta z\lambda(k_x^2 + k_y^2)]\exp[-2\pi i(k_x x' + k_y y')]dk_x dk_y\right|^2 \quad (2.4)$$

where λ is the wavelength of the electrons and *Δz* is the defocus. This equation can be solved numerically for different specimen geometries with a given magnetization.

An illuminating aspect of Equation 2.4 is that the contrast is adjusted by the defocus *Δz*. This means that an important part of setting up the microscope in Lorentz mode is to actually intentionally operate it out-of-focus. There is an optimal defocus to get the strong contrast asymmetry in Figures 2.28 and 2.29 and this amount of defocus will depend on the size of the structures being imaged and the microscope magnification. Aside from finding the correct defocus, the rest of the microscope alignment for Lorentz imaging is generally straightforward. The rest of this paragraph is intended to help someone who is familiar with the standard operation of a transmission electron microscope to achieve good Lorentz contrast images. I tended to use the largest condenser aperture to be able to image large crystals and large sample areas. I always wanted the brightest possible beam so I chose spot size 1 and made sure to align the gun so that the spot was the brightest it could

be. If this gave a brightness asymmetry across the image I would then slightly adjust the gun or the beam itself so the whole crystal was evenly illuminated. I always used convergence angle setting $\alpha = 3$. For large crystals I used a magnification typically 400-600x and for the small kagome crystals I used around 1000-1200x, always in low-mag mode. I would adjust the condenser stigmation as usual and try to have a fairly widely spread beam so that the phase of the electrons was constant across the crystal. The beam was about 1/3 to 1/2 the diameter of the fluorescent screen. The one other main adjustment to improve the Lorentz contrast asymmetry is to change the intermediate lens (IL) stigmation if the different ASI sublattices have different amounts of asymmetry or different Lorentz contrast. It can be difficult to get all sublattices to look the same, especially in the kagome geometry. A lot of the IL stigmation adjusting is trial-and-error. I always capture images with CCD bin size 1, typically for 3 seconds for large crystals. Bin size 1 exposures for 1.5-2 seconds in search mode are helpful in adjusting the defocus and IL stigmation.

The small contrast asymmetry across each element shown in Figure 2.28 is all that is necessary to determine the magnetization direction of that element. Despite the challenges of working with the SiN TEM membrane substrates, I developed my fabrication process for $FePd_3$ ASI samples on top of the membranes in order to utilize the great advantages of the Lorentz TEM for imaging of the ASI magnetic configurations. Lorentz TEM is fast. It only takes about 3 seconds of exposure on a large crystal to capture a high quality Lorentz TEM image. For the smaller crystals found in Section 3.2, exposures were only 0.5 or 1 seconds. For the large crystals, this is about 200-400 times faster than MFM for a similarly sized crystal, which can take up to 10 or 20 minutes per image.

Besides MFM, PEEM is the next most common imaging technique for ASI structures. It is fairly fast, with individual exposures on the order of 12 seconds, but requires highly specialized x-rays and predetermined beam time at a cyclotron facility. Users have limited imaging during their allotted time and the tool is not readily available. In contrast, the TEM it is significantly faster and it is easier to move the sample around to quickly look at different areas and either find the best crystals to image or take many images in one session. In general, Lorentz TEM has huge advantages over other techniques because it allows for the collection of significantly more data. I can image more crystals, more copies of similar geometry crystals, and the same crystals multiple times to get better averages and more precise measurements of ordering behavior, representative of true properties and statistics. The tool is readily available at UMD for frequent and versatile use. Lorentz TEM allowed Daunheimer to collect enough data with small enough error bars to fit a power law behavior for avalanche sizes in the kagome geometry shown in Figure 1.31 and it has enabled me to gather thousands of images that will be summarized and discussed throughout Chapters 3 and 4.

Lorentz TEM allows for unambiguous imaging of connected kagome structures. Connecting kagome elements increases their magnetostatic coupling and has also been shown to reduce disorder in the samples [94], ideal for realizing the theoretically predicted long-range ordered states in thermally active ASI. Lorentz TEM is sensitive to the in-plane component of the magnetization as opposed to MFM, which is only sensitive to the out-of-plane component. In disconnected structures, MFM can uniquely determine the exact spin configuration (shown, for example, in Figure 1.34), but if the elements are connected, there

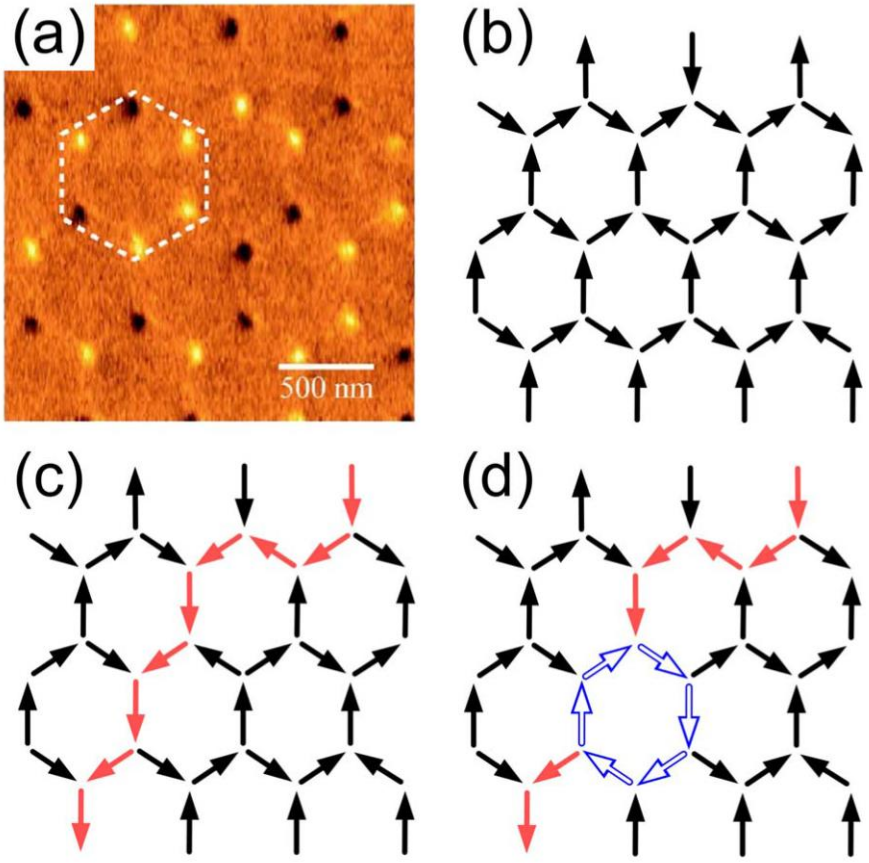

Figure 2.30 **Ambiguity in using MFM to image connected kagome structures**. **a.** MFM image of a connected kagome structure and corresponding proposed spin map (**b**) from Tanaka *et al*. [124]. **c, d.** Alternative valid spin maps created by reversing chains of spins, shown in different colors. Figure adapted from Refs. [93] and [124].

is ambiguity in MFM imaging of the kagome lattice. This ambiguity is demonstrated in Figure 2.30.

2.3.3 Lorentz image processing

Often times in experiments I capture hundreds of Lorentz TEM images, each with thousands of individual magnetic elements. It is impossible to go through by hand and assign a direction to each element, so I use MATLAB scripts to process the images (see

Appendix in Ref. [132]). The general outline of the scripts is to first find the coordinates of the individual elements, then go through one-by-one to analyze the asymmetry in each element and assign a direction to it. Then the script can output a variety of data to help understand the magnetic configurations. The scripts can calculate ordering parameters that will be discussed later, or simply output a text file with all of the information about spin directions and locations that can later be used to do more rigorous statistical analysis. The scripts can also output visual representations of the data to aid in interpreting and analyzing the configurations. This could mean printing colored lines or arrows on top of the Lorentz images to make it easier to see the individual element directions, putting labels or dots on the vertices to illustrate what type of vertex it is in the square geometry or what the charge of the vertex is in the kagome geometry, or overlaying color to represent the domains of ordering observed in the kagome geometry. All of these outputs will be shown later, for example, in Figures 2.35, 3.15, and 3.26.

Stephen Daunheimer wrote the main analysis script, named the "crawler," for the kagome geometry along with help from our advisor, John Cumings. I modified this script to be able to analyze the small kagome crystals found in Section 3.2 and also added vertex plotting, charge domain plotting, and other analysis to investigate charge-ordering. Dr. Cumings wrote a different crawler to analyze the square geometry samples which I also modified to plot arrows on the Lorentz images. I worked for one summer with a high school student, Thomas Marsh, who adapted the existing crawlers to suit the topological defect samples discussed in Chapter 4. Thomas was very successful at making the crawler highly versatile and adaptable to any geometry we wanted to study and also implemented the vertex notation found in Figures 2.36 and 4.51 as well as many other figures in Chapter 4.

I can now explain why the bright spots on the SiN membranes discussed in Section 2.2.5 were so detrimental to data collection. The crawler is extremely sensitive to the slight asymmetry in the contrast of individual elements in the Lorentz TEM images. If the crawler is not entirely sure of the moment direction, there is a mode where it allows the user to double check or decide the doubted elements in question. This makes it possible to correct some data that may not have as high an asymmetry as desired. The problem with the really bright spots is that they throw off the contrast asymmetry so much that instead of just being unsure about an element and allowing for user input, the crawler is sure that the element is pointing in the direction opposite from what it actually is and the crawler simply moves on. In normal samples without spots, it is still possible for the crawler to get a few elements wrong and not doubt them so I make sure to double check all of the elements in important images submitted for publication, included in conference presentations, and found in this dissertation.

The quality of the $FePd_3$ thin film also affects how well the crawler is able to process the images. If the film does not have a high magnetization, there will be little or no asymmetry in each element and the crawler will doubt most of the spins and get more of them wrong. It can be very difficult to see any symmetry sometimes. Films with higher magnetization result in better ordering, but the actual quality of the Lorentz data is better too because there is more asymmetry. I try to select films with the best Lorentz contrast before patterning because these will directly result in better data and more accurate analysis once patterned.

2.4 Heating experiments

2.4.1 Annealing protocols

In order to see thermal effects in my ASI samples, I need to heat the samples. I investigated a wide variety of annealing protocols throughout many different experiments. Samples were either heated *in-situ*, inside the microscope with the heating holder, or *ex-situ* in an Ar atmosphere inside a tube furnace. The heating holder was beneficial during early work and some of the first heating experiments because it provided immediate information on whether any elements had flipped. I could polarize a sample along one axis using our custom electronics, heat it, and immediately tell whether the temperature I had heated to was hot enough to thermalize the system and flip any elements. After some initial checks it became less useful to use the heating holder on the patterned samples. It was still good for checking the T_C of thin films, but it heats the membrane very asymmetrically. The heating holder clamps the Si frame of the substrate around the outside and heat is transferred through the clamp, meaning that all the heat comes from the edges of the substrate. The thin SiN membrane radiates heat from the middle, so as the edges heat up, the middle of the membrane stays considerably cooler than the outside. This is shown schematically in Figure 2.31. There is a strong temperature gradient across the membrane with different parts of the membrane at different temperatures.

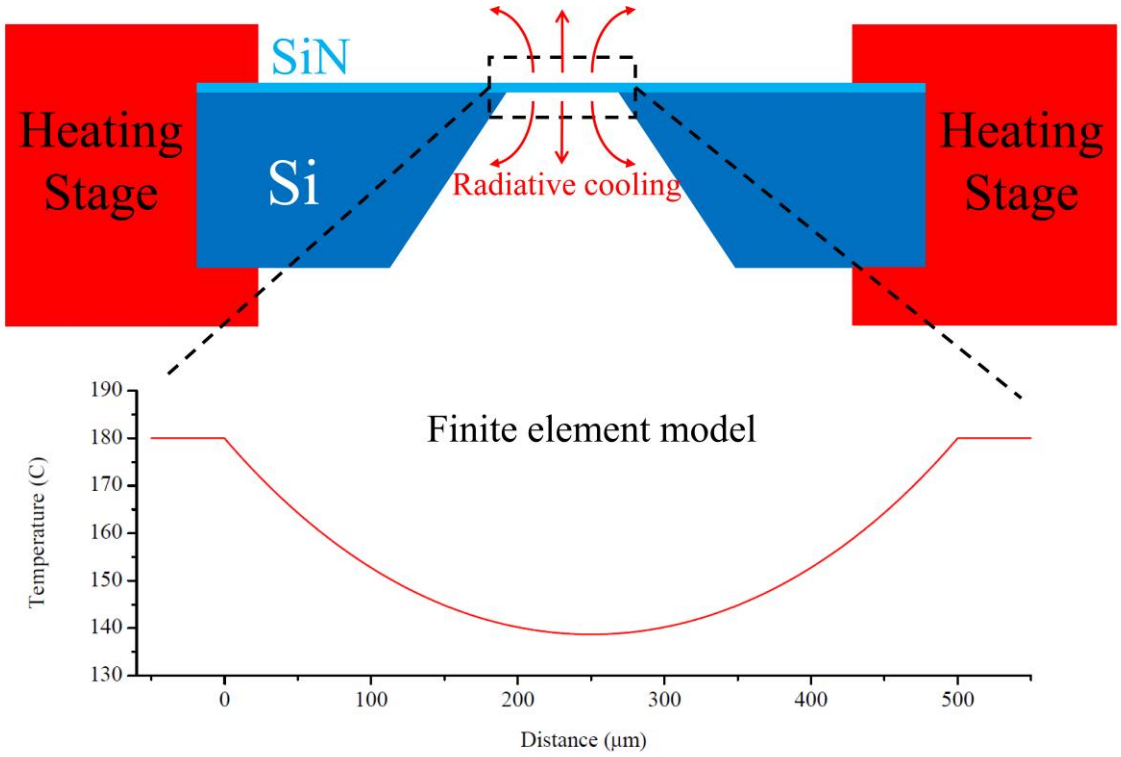

Figure 2.31 **Heating holder radiative heat loss**. Schematic of the *in-situ* heating holder. The holder clamps onto the SiN substrates and provides heat from the edges. The thicker Si conducts heat well, but there is a lot of heat lost to radiative cooling from the middle of the thin SiN membrane. A finite element model temperature vs. position profile shows how the temperature drops quadratically from the edges of the membrane held at constant temperature. The middle of the membrane is about 40 °C colder than the edge, creating a very strong temperature gradient across the membrane. Cumings' group graduate student Hanna Nilsson helped me with the finite element model.

The heating holder has an external controller where the temperature can be set at will. We typically take this temperature to be the temperature at the edge of the SiN membrane because there is much less radiative cooling from the thicker Si frame. However, as can be seen in Figure 2.31, there is approximately a 40 °C temperature gradient between the edge of the SiN membrane and the middle. Also, even the temperature at the edge may not be very accurate as well. This makes it difficult to examine specific crystals because they all have different temperatures where they become thermally active and the real temperature at the crystal is not well known. The heating holder also heats and cools rather

Figure 2.32 **Tube furnace**. The samples are inserted into the glass tube and positioned in the middle of the furnace. Ar gas flows from the cylinders on the right through the furnace to heat the samples without oxidizing them. To quench the samples, the glass tube can be quickly slid far to the right so that the samples leave the middle section where the heat source is.

slowly and uses valuable TEM time to do the heating. It requires a lot of manual operation to set and ramp temperatures and is also hard to control the temperature overall. After some initial testing with the heating holder where thermal flipping of individual magnets was observed, I moved towards using a tube furnace to heat my samples. A photograph of the tube furnace is shown in Figure 2.32.

The furnace has a long glass tube that sticks out on both ends from a middle section where the heating elements are. Ar gas flows in from a cylinder and flow controller on one side and flows out into a beaker filled with oil in the fume hood on the other side. I load my samples onto a quartz boat covered in aluminum foil and slide the boat into the middle of the furnace. First I flood the tube with Ar gas to get as much oxygen out as possible, then reduce the flow to a very slow rate. I calculate the flow rate so that the Ar will spend

3-5 minutes passing through the furnace, giving it enough time to heat up and efficiently heat the samples inside.

A typical tube furnace annealing protocol I used is shown in Figure 2.33. Most samples went through this protocol, where they are heated in Ar atmosphere above their T_C and cooled at a rate of 1 °C min^{-1} back down to room temperature. When the sample is above T_C, the elements have no magnetization. As they are cooled below T_C, they regain a random magnetization and are then able to interact and flip to arrange themselves in an energetically favorable configuration. The magnets continue this until they reach their blocking temperature (T_b) when they stop flipping. I cool the samples at the given rate until they are safely below T_b and then can cool faster to save time. The resulting magnetic configurations can then be imaged using Lorentz TEM.

The annealing protocol in Figure 2.33 is typically accomplished by simply turning off the tube furnace and letting it cool down on its own. The cooling may not be strictly linear, but I have observed the cooling rate and it is close. I tried many other annealing protocols as well, especially in early experiments. These include heating samples to a given temperature, above or below T_C, holding there for anywhere from 5 minutes to 10 hours, and then quenching the sample. Quenching typically involves quickly sliding the glass tube far to one side, so that the boat with samples in it would move out of the middle of the furnace where the heat is coming from. This, in addition to rapidly turning up the Ar flow, helps cool the samples as fast as possible. I also tried very slow cooling, reducing the temperature by 1 °C every 5 minutes, 20 minutes, 1 hour, or 5 hours. In general, I

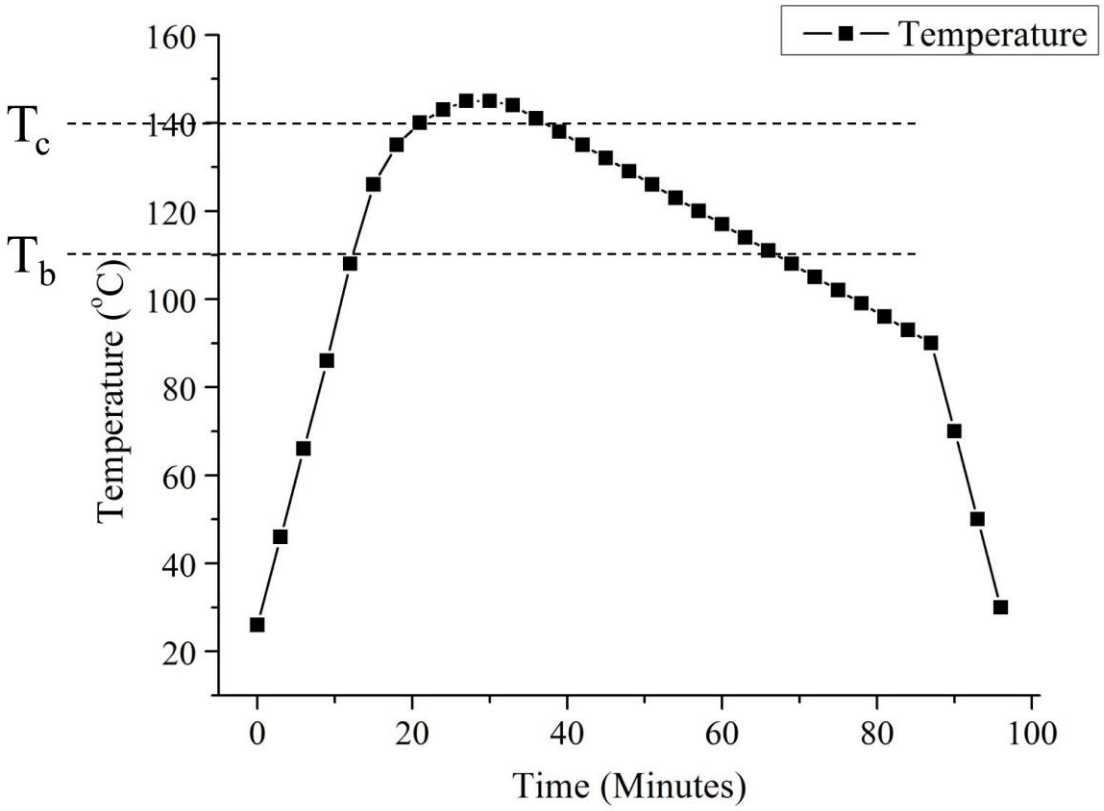

Figure 2.33 **Annealing protocol**. Typical annealing protocol for most data in this dissertation. Samples are first heated quickly above their T_C and then cooled slowly at 1 °C min^{-1} back down to room temperature. The magnets are able to flip when the temperature is above T_b and so they are able to flip and interact in the region between T_C and T_b on the graph. The T_C and T_b indicated are for more recent samples and can vary slightly depending on the deposition of $FePd_3$.

found that faster cooling tends to give more ordered states. The longer anneals may work better if the temperature could be more precisely controlled. The controller on the furnace has limited programmability, and it typically oscillates by 1 °C when it should be holding at a constant temperature. This does not result in linear cooling. Longer anneals apparently give elements more time to flip into non-ordered states and get stuck there. The process of heating above T_C, cooling at 1 °C min^{-1}, and quenching tends to give the best results. Most of the data in later chapters comes from an annealing process like that in Figure 2.33. Any

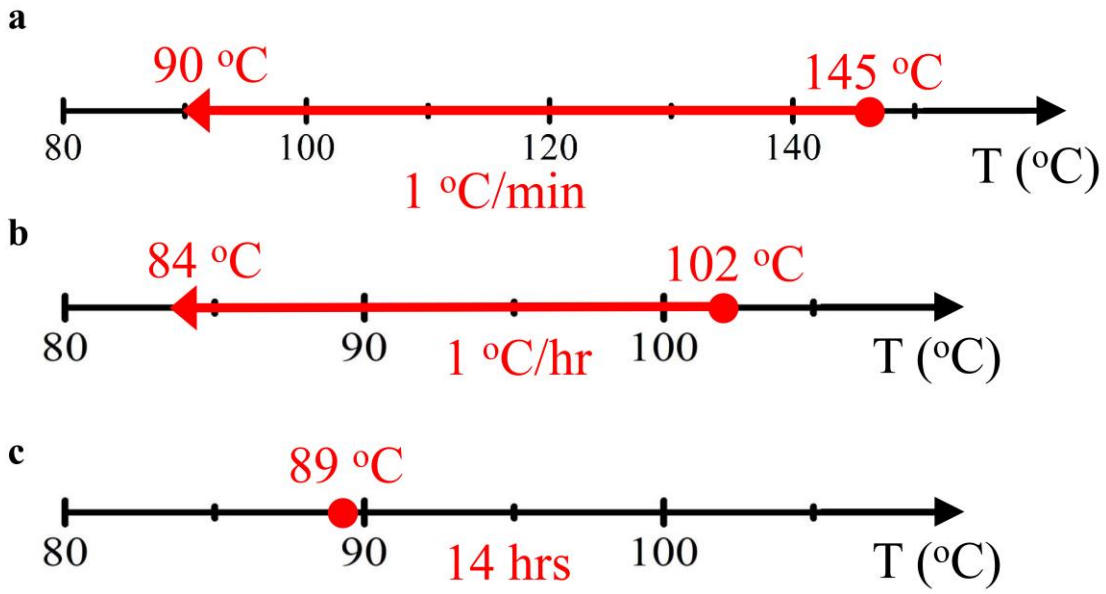

Figure 2.34 **Other annealing graphs. a.** Alternative, compact representation of the annealing process in Figure 2.33. Samples are heated to 145 °C and cooled at a rate of 1 °C min^{-1} down to 90 °C. **b**. Different annealing process for most small kagome crystals with lower T_C. Samples are still heated above T_C and cooled at a given rate down past the temperature where spins stop flipping. **c.** Quenching schematic. Samples were heated at 89 °C for 14 hours and then cooled as fast as possible.

image or data from a different process will be noted with a graph similar to Figure 2.34, but otherwise it should be assumed the data came from the typical process in Figure 2.33.

After annealing, the samples are imaged in the transmission electron microscope. Before inserting samples into the microscope, a degaussing procedure is run on the microscope's objective lens to remove any remnant field (< 1 gauss) at the sample that could influence the magnetic configuration. This is done through custom electronics to take control of the objective lens as well as a LabVIEW program to run the degaussing procedure. The program to degauss the lens was written by Stephen Daunheimer and is similar to Protocol 1 in Figure 1.20. The degaussing procedure is necessary because the lens is capable of producing extremely high fields and in order to do so, it has a soft

magnetic core that amplifies the field the solenoid generates. This core will have a non-zero magnetization when the solenoid is turned off which can affect the magnetic configuration of annealed ASI samples and bias the data. To check that we are at < 1 gauss, I use a special TEM holder with a Hall bar sensor mounted on the end. The sensor measures the magnetic field at the exact location the ASI sample is going to be. This is important because the specific LabVIEW input parameters needed to degauss the lens change over time, so the success of the degaussing must be verified and the parameters tweaked until the Hall sensor measures < 1 gauss at the sample. I try to run the degaussing procedure with all other settings on the microscope close to what they will be during Lorentz imaging. This insures that stray fields from other lenses used in Lorentz imaging will be cancelled during the degaussing and the field will remain close to zero at the sample during imaging. I have tested the stability of the setup and find that after degaussing, the field at the sample remained at < 1 gauss for at least 7 hours.

2.4.2 Initial results

The foundations of the project and this dissertation are built upon thermal activation of ASI using a new material with strong experimental advantages. An important benchmark test for the $FePd_3$ ASI system I developed is to see whether the square ice ground state can be observed after annealing. This ground state was previously reported in both as-grown samples [77] and in permalloy samples carefully heated above their T_C [78,82]. In order to verify that $FePd_3$ is a good material for these studies, the square ground state should readily be reproduced. I found that upon heating, I observe thermally activated reversal of individual magnets when the $FePd_3$ ASI samples are heated to about 110 °C, though in

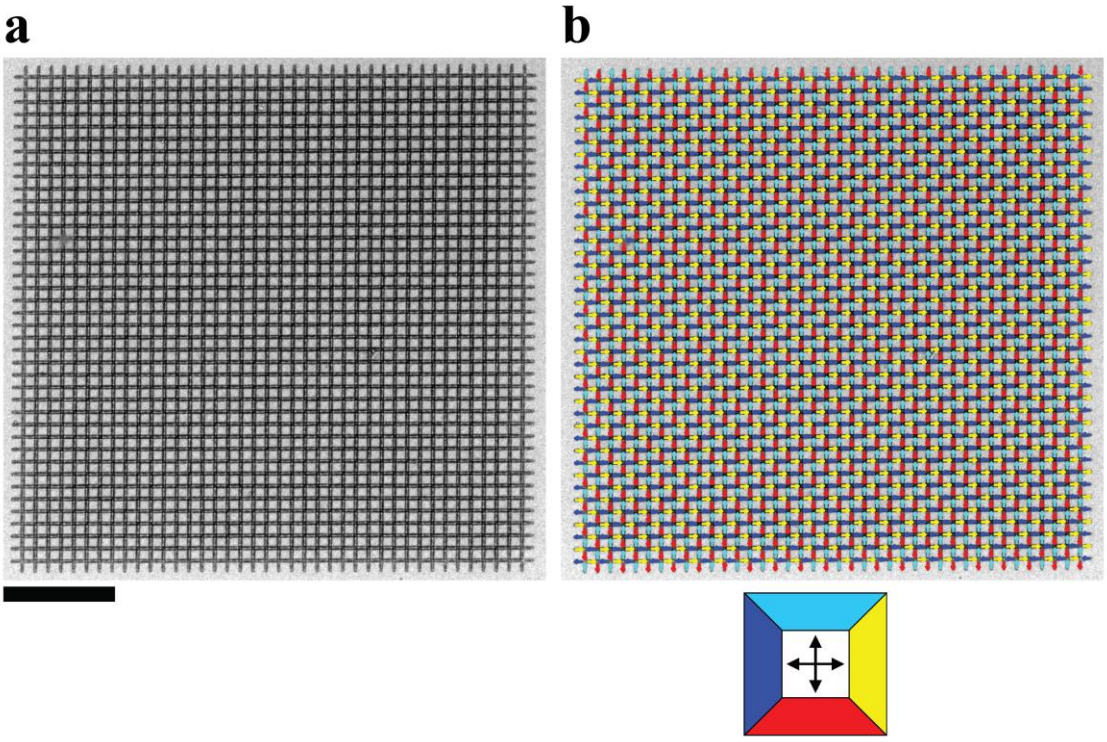

Figure 2.35 **Square ice ground state**. Lorentz contrast TEM image (**a**) and corresponding spin map (**b**) of a square geometry sample that was heated above its T_C and cooled back to room temperature showing perfect ground-state ordering of type I vertices. Scale bar is 4 μm.

some early samples thermal activation began around 88 °C due to different $FePd_3$ film properties. After the annealing protocol in Figure 2.33, I observe the square geometry samples to readily find the ground state, often with perfect or nearly perfect ordering of type I vertices. The largest perfect ground-state domains we have observed are 30 μm × 30 μm, limited only by the size of the samples studied. An example of a perfectly ordered square sample is shown in Figure 2.35.

Figure 2.35 shows a Lorentz contrast image of a square geometry sample that was annealed as well as a corresponding spin map added by the crawler described in Section 2.3.3. In some crystals and especially in the crystals discussed in Chapter 4, it is beneficial to switch to a vertex notation for the square ice. This notation is shown in Figure 2.36

where a colored dot is added to each vertex to signify its energy level. For perfectly ordered crystals the new notation is not very enlightening, but in some crystals where domains of ground-state order are separated by domain walls of higher energy vertices, the vertex notation is very helpful at illuminating the behavior of the samples. An example of this is shown in Figures 2.37 and 2.38.

To understand Figures 2.37 and 2.38, I note that the square ice ground state is two-fold degenerate, corresponding to the two possible tilings of alternating type I vertices. Therefore, there can be different domains of ground-state order separated by a domain wall. This is found in some square geometry samples, though they still tend to be very well ordered and support large domains of ordering. We can quantify the degree of ordering in the square geometry with a correlation parameter I will call "*C*." *C* is found by assigning a value of 1 to a pair of neighboring spins of they are in a head-to-tail arrangement and a -1 to the pair if they are head-to-head or tail-to-tail arrangement, adding up the contribution from all pairs, and then dividing by the total number of pairs. A *C* value of 1 represents perfect ordering, 0 is disordered, and -1 is anti-ordering of all type IV vertices. The crystal in Figure 2.35 has $C = 1$, Figure 2.37 has $C = 0.95$, and Figure 2.38 has $C = 0.92$.

I note that I fabricated square geometry samples with both three-fold edges (3FE) — shown in Figure 2.38 — and one-fold edges (1FE) — shown in Figure 2.37. 1FE crystals have elements sticking out at the edges and 3FE crystals have flat edges. We observe no significant differences in the degrees of ordering between the square geometry crystals with different types of edges. Overall, I captured over 200 Lorentz TEM images of different annealed square geometry crystals fabricated from 5 different samples made from different depositions of $FePd_3$. Every single crystal has a *C* value greater than 0.85, with most 0.95

or higher. These are incredibly robust results and demonstrate the reproducible and well-understood behavior of long-range ordering in square ASI.

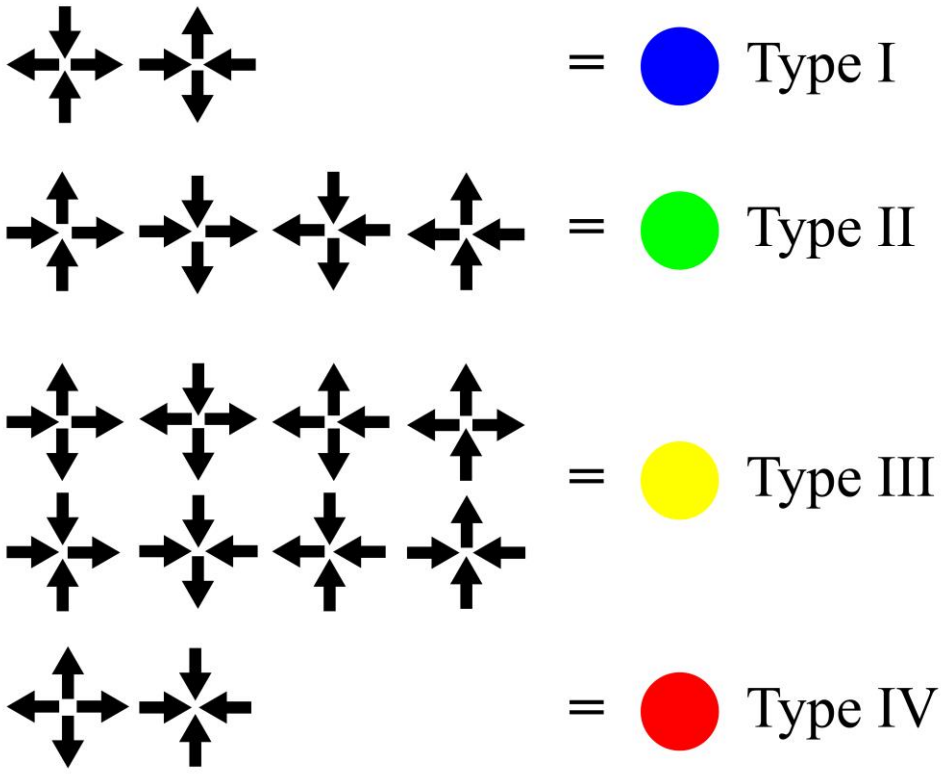

Figure 2.36 **Vertex types notation**. Definitions of the 16 square geometry vertex types grouped by energy level. The colored dots are placed on top of the vertices in Lorentz TEM images to aid in data analysis and visualization. This vertex notation is used in Figures 2.37 and 2.38 as well as extensively throughout Chapter 4.

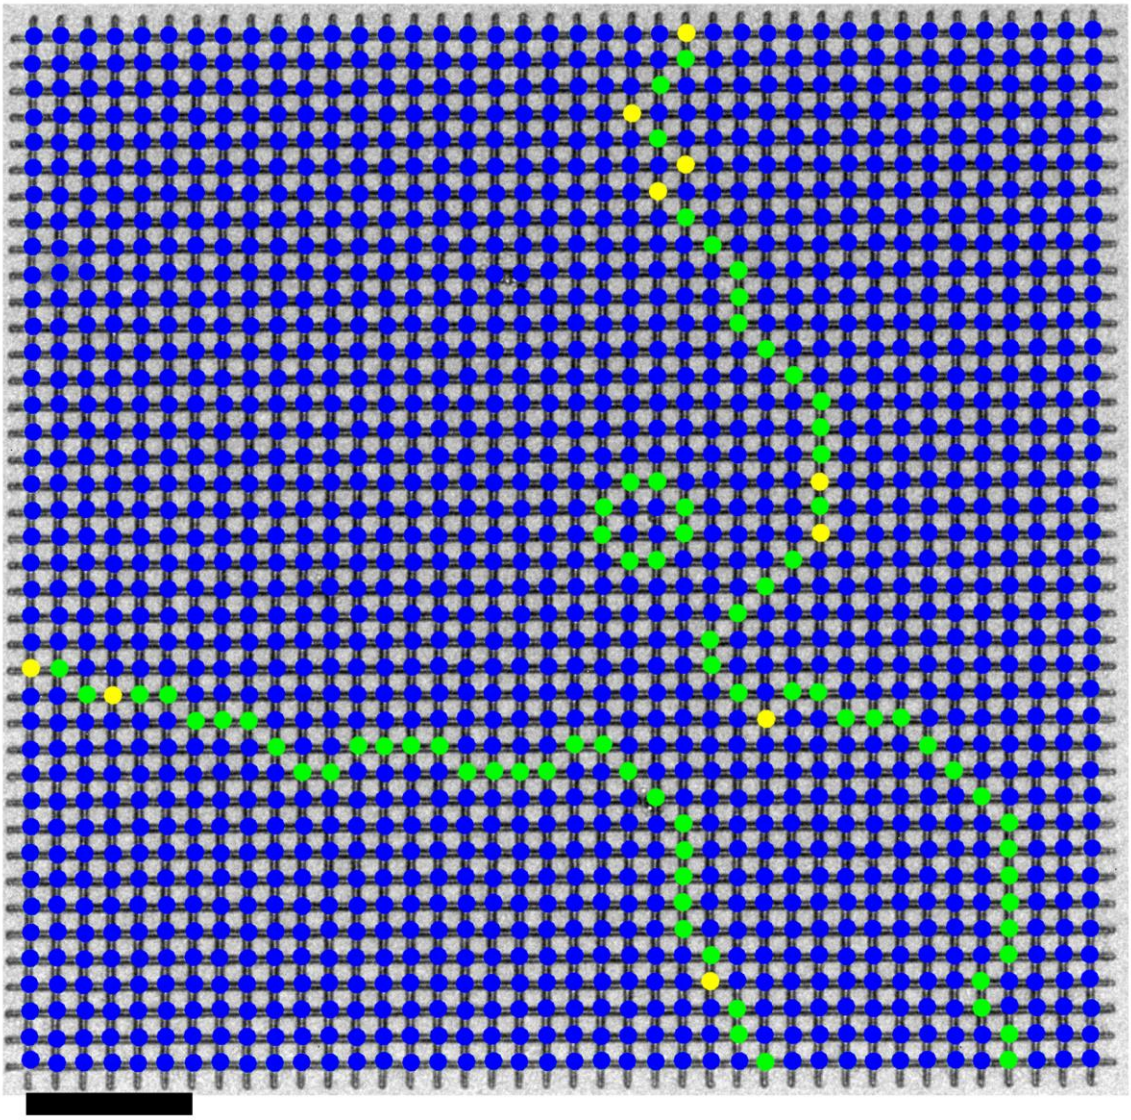

Figure 2.37 **Square ice ground-state domains**. Vertex notation added to a Lorentz TEM image of a square geometry crystal with 1FE that has been annealed by the process in Figure 2.33. The crystal shows large regions of ground-state ordering of type I vertices separated by domain walls made of higher energy vertices. There is one isolated domain as well as two domain walls that extend to the edges of the finite crystal. Domain walls are typically made of type II and III vertices, consistent with other experimental and theoretical studies of thermalized square ASI [78,119]. We note that we have never observed a type IV vertex in any of our samples. The correlation parameter *C* for this crystal is 0.95, showing very good ordering. Scale bar is 3 μm.

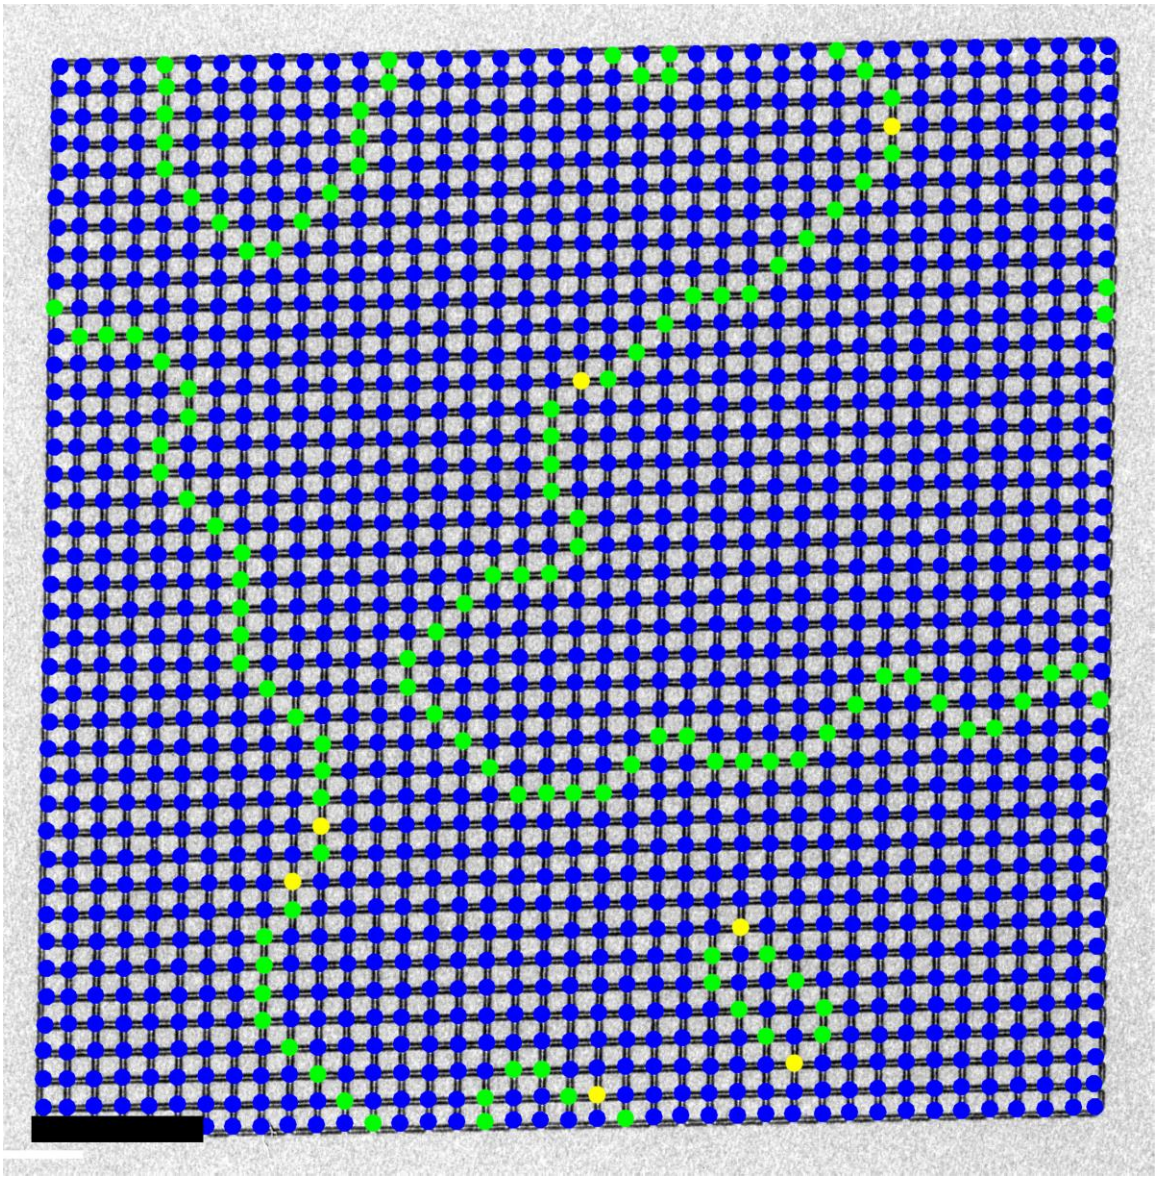

Figure 2.38 **Square ice domains**. Vertex notation Lorentz TEM image of a square geometry crystal with 3FE. C = 0.92 for this crystal. Scale bar is 3 μm.

The consistent observation of square ice ground-state ordering confirms $FePd_3$ is a good material for studying thermally active ASI systems and it reproduces previous results, though the $FePd_3$ has some major advantages over other groups. Zhang *et al*. report a very small temperature window where they can conduct experiments with permalloy samples [78]. They see magnets begin to flip at 545 °C, but if their samples are heated 560 °C, they

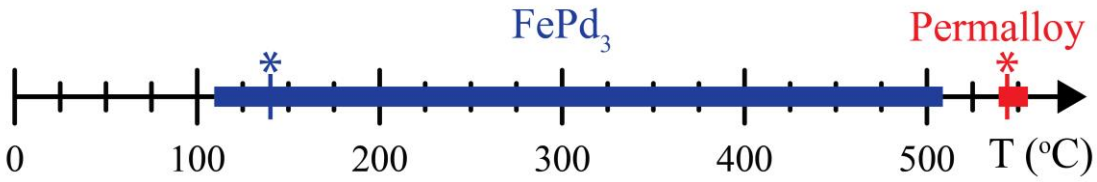

Figure 2.39 **Experimental temperature window**. Experimental temperature window comparing $FePd_3$ and 25 nm thick permalloy. Lower bound is the temperature where magnets begin to flip due to thermal activation, higher bound is limited by nanostructure degradation. Stars indicate the relative Curie temperatures of the two materials. Permalloy data reported in Ref. [78].

begin to see structural degradation of the nanomagnets due to lateral diffusion. They also report having to specially treat the substrates of their samples to obtain even this small experimental temperature window. A small window makes experiments very challenging and makes repeated experiments extremely difficult as the samples cannot last very long. We have tested our samples by heating them up to 510 °C and did not detect any structural degradation at that temperature. Thus the $FePd_3$ ASI system gives a much wider temperature window compared to permalloy and allows for greater flexibility and ease in performing experiments. A comparison of the experimental temperature windows for the two materials is shown in Figure 2.39.

As noted above, I found that the perfectly ordered square ice ground state was achieved relatively easily with $FePd_3$. The behavior was extremely robust and can be observed quite quickly as well. The full annealing protocol in Figure 2.33 takes around 90 minutes to complete, but the square ground state can also be observed with quenching and in that case can only take 20-30 minutes to complete the heating. This is significantly faster than the perfect square ordering reported by Farhan *et al*. which took around 8 hours to achieve [82]. In their case, they only image one crystal, while in my work I can heat one

sample for 30 minutes, quench it, and capture images of tens or hundreds of crystals in less time than it takes them to observe one. Zhang *et al*. also do not report any perfectly ordered square crystals like the one in Figure 2.35. The ease of experiments and robust ground-state ordering in $FePd_3$ ASI are huge advantages that enable the great flexibility in sample design and the plethora of data gathered during my dissertation research. This data and all experiments performed in addition to the ones in this section will be presented in Chapters 3 and 4.

2.5 Numerical methods

2.5.1 Introduction and general framework

Numerical modeling and simulation are essential tools in virtually all areas of physics. Simulations are useful for making predictions about experimental systems as well as give new insights on observed experimental behavior and properties. All of the methods discussed thus far in this chapter have been experimental methods for fabricating and investigating thermally active ASI samples. This section will provide an introduction to and the general framework for the computational methods that I use in combination with these experimental methods to better understand the behavior and long-range ordering properties of my $FePd_3$ ASI samples.

I use a Monte Carlo method, as Monte Carlo simulations have proven to be a successful approach in understanding complex behavior in other ASI systems [83,119,141,183,184]. I will describe the Monte Carlo technique that I employ and give the details of how the method is implemented. Often, theoretical studies and numerical

methods predict an experimental behavior, but upon investigation the experimental behavior deviates from the prediction. In this instance it can be beneficial to go back to the numerical model to help understand what may be causing the deviation, modify or improve the theoretical model and understanding, and make new predictions. There have been many predictions of long-range order in artificially frustrated systems which motivated some of my studies, especially in the kagome geometry [51,78,95,98,125]. My numerical modeling provides some insight as to why the ordered states have proved difficult to observe.

Monte Carlo simulations are very common and very powerful computational methods that span the physical sciences [185,186]. There are many different branches and methods that use different algorithms suited for different applications. The Metropolis-Hastings algorithm is one common method that is very versatile [187–189]. It can be used on a variety of complex systems, but as applied to ASI studies, the basic premise is as follows: First, randomly select a spin in a lattice and calculate the change in total energy of the system if that spin were to flip. If flipping the spin will lower the total energy of the system, flip the spin. If flipping the spin will raise the total energy of the system, flip the spin with a probability that depends on the temperature of the system. Higher temperatures are more likely to allow for excitations that raise the total energy of the system. Some thermal ASI studies very successfully use a Metropolis algorithm to model observed behavior [78,85].

I found a kinetic Monte Carlo technique [190] to be the most useful in modeling my $FePd_3$ ASI samples. The kinetic Monte Carlo technique was originally developed to model vacancy migration in a binary alloy, but it has also been found to be an accurate method for describing thermally active ASI systems [81,82,95]. The first step is to create

a lattice of the ASI spins. I choose the geometry, size, and boundary conditions of the lattice to match those of my experimental samples. Then, an Arrhenius flip rate is assigned to each element given by [142,190]

$$\tau^{-1} = \nu_0 \exp(\frac{E_0 + \Delta E}{k_B T}) \tag{2.5}$$

where E_0 is an intrinsic energy barrier, T is the temperature, v_0 is a material specific prefactor, k_B is Boltzmann's constant, and ΔE is half the change in total energy of the system if that spin were to flip. ΔE can be calculated from magnetic Coulomb interactions or dipolar interactions depending on the geometry and the system being studied. E_0 is known to be well-approximated by the shape anisotropy energy of the magnetic element and was given in Equation 1.14 [81,87,139,145]. I always use M = 200 kA m^{-1}, T = 420 K, and $v_0 = 10^{-25}$ s^{-1}. E_0 depends on the demagnetizing factor D_Z which is a purely geometric term derived from the shape of the magnet and how easy it is to flip the magnetization along different axes. D_Z is the factor along the long axis of the magnet and has been solved numerically for a general ellipsoid [25] and exactly for a rectangular prism [145]. We approximate the ASI nanomagnet as a rectangular prism for simplicity in the calculation and we believe this is a reasonable and good approximation. We are less concerned with the values of parameters T and v_0 as these are generally used to evolve time in the model and affect every rate in the lattice in the same manner. We are more interested in the long-time, steady-state ordering in the system and observing how ordered the lattice can become.

The magnets start in a random initial condition. The flip rates in Equation 2.5 are calculated for each spin and all the rates are summed. The rate for each spin is then divided by the total, yielding a relative probability distribution for the individual magnets to flip. One magnet is then selected at random, weighted by the probability distribution, and

flipped. The rates are recalculated and the process repeats. A visualization of the state of the individual magnets is given in Figure 2.40. The magnet has two possible states in an effective double well potential. There is some intrinsic energy barrier E_0 to overcome in order to flip and a change in energy ΔE which can be positive or negative if the magnet flips. Both E_0 and ΔE are typically on the order of a few eV for our experimental samples and simulation parameters. The benefit of the kinetic Monte Carlo technique is that spins that are in less favorable energetic arrangements are more likely to flip and spins that are in energetically favorable environments are unlikely to flip, though it is still possible. This is characteristic of real samples finding ordered states. Well-ordered regions are unlikely to thermally excite and raise the total energy, though it is still possible. This technique also reduces the time it takes to run the simulations. There is little processor time wasted randomly picking spins that are already in low energy states. The model picks the spins that should flip first, quickly achieves decent ordering, and then the steady-state behavior progresses.

Disorder is known to play an important role in the long-range ordering of ASI systems and it is important to be able to incorporate disorder into the simulations. I discussed in Chapter 1 how different types of disorder have similar effects on the energy landscape of the ASI vertices (see Figure 1.27). Therefore, it is essential to include the effects of disorder in the system, but less important exactly how the disorder enters the model. As we will see in later chapters, there is a characteristic spread in the widths of individual magnetic elements due to artifacts from the electron beam lithography. This is

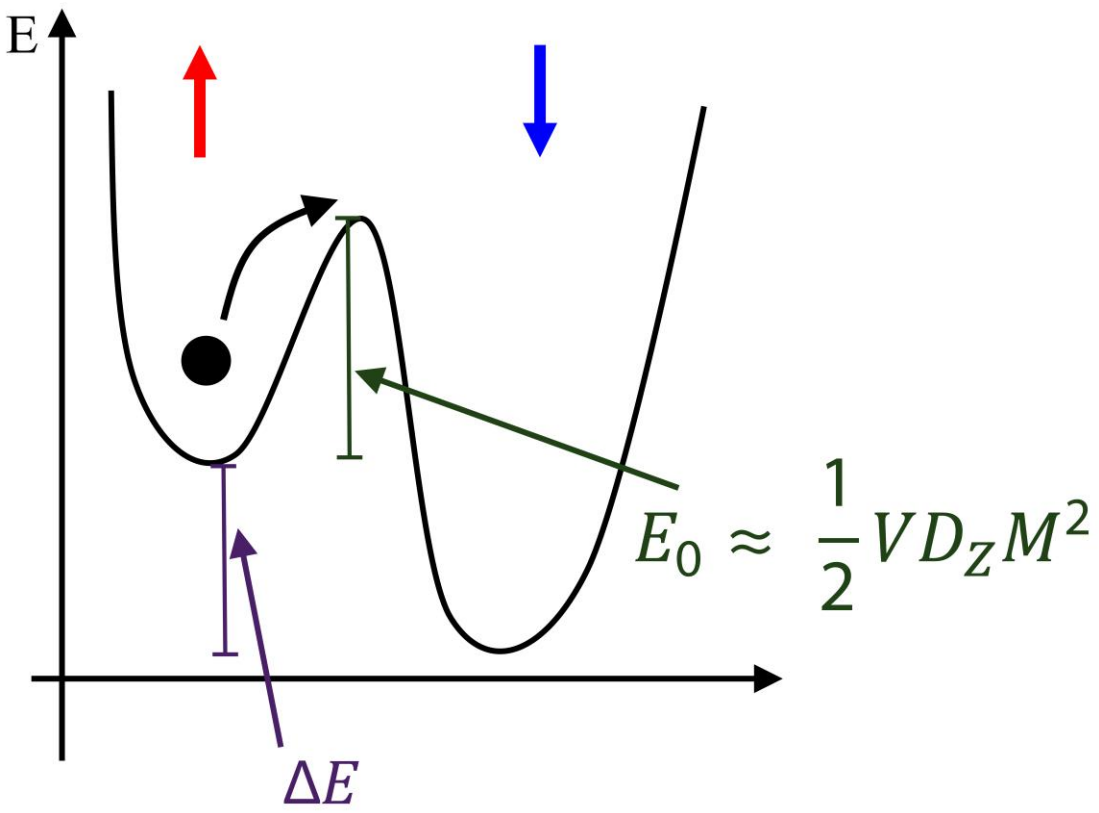

Figure 2.40 **Spin state double well potential**. Schematic of the energy landscape for one ASI nanomagnet. There is an intrinsic energy barrier E_0 to overcome in order to flip which is approximated by the shape anisotropy energy of the element. ΔE is the change in energy if the magnet were to flip. Depending on the current state of the magnet, ΔE can be positive or negative. E_0 is always positive.

a realistic source of disorder in experimental samples and is relatively easy to implement in the simulations. Individual spins in the simulation lattice are randomly assigned a width sampled from a Gaussian distribution with a characteristic mean width and standard deviation (σ_w). The given mean width and σ_w depend on the exact system and geometry being studied, but are typically on the order of 120 nm and 1 nm, respectively.

The effects of this width disorder enters the simulations in two ways. First, the intrinsic energy barrier E_0 depends on both the volume of the element and D_Z which are both strong functions of the width of the magnet. This leads to a range of E_0 values assigned

to each spin, heavily affected by the actual width of the magnet. Second, the ΔE term will be a function of the magnetic moment of the individual magnet. ΔE can be calculated in different ways, as will be discussed in Chapters 3 and 4, but it will always depend on the magnetic moment, which is the product of the nanomagnet's magnetization and volume. This will cause slight variations in the interaction energies across the sample. Both E_0 and ΔE have a strong dependence on the widths of the magnetic elements which propagate the disorder into the energy landscape of the vertex interactions and spin flips. This can cause certain spins to be very favorable to flip, and others to be very unfavorable, potentially pinning spins and defects that break the long-range order in the system.

2.5.2 Square geometry Monte Carlo

We observe the experimental square geometry samples to readily find large domains of ground-state order after annealing. It is useful to simulate these experiments to gain greater understanding of the system. The simulations follow the basic framework laid out in the previous section. The square geometry samples are modeled by elements 120 nm wide, 23 nm thick, and 400 nm long and the simulations are the same size as experimental samples with 40 elements per side. We use dipolar interactions for the coupling between magnets to calculate ΔE, similar to other reports [82]. The formula for dipolar coupling is found in Equation 1.11. The magnetic moments are given by a unit vector pointing along the length of the magnet, the magnetization of the material 200 kA m^{-1}, and the volume of the element is nominally 1.10×10^{-21} m^3. The simulations start in a random configuration and we evolve them for an average of 200 flips per spin. We investigate the effects of

different degrees of width disorder σ_w as well as different numbers of nearest neighbors to include in the dipolar interactions. Dipolar interactions are typically long-range, but including interactions between one element and all other elements in the lattice is computationally expensive, therefore it is common to truncate the interactions at some distance. We find that including only nearest neighbor interactions does an excellent job of reproducing our experimental data and including higher order interactions results in qualitatively similar results, justifying truncating the interactions for efficiency.

With no disorder the simulations produce perfect ground-state ordering every time. A perfectly ordered final state from a simulation is shown in Figure 2.41. All the simulation images use the vertex notation defined in Figure 2.36. When including disorder, we generally repeat simulations a number of times with different random sampling of disorder because a different sample can yield different results and domain configurations. Some simulations show a perfectly ordered lattice with as much at σ_w = 1 nm disorder, though this is not generally the case. Figure 2.42 shows 4 typical domain patterns observed in simulations, produced with varying width disorder from σ_w = 0.5-2.5 nm. This level of disorder captures the experimentally observed large domains of type I vertices and it also reflects a realistic spread in the actual widths of the lithographically patterned arrays as observed with high-resolution TEM. We can calculate the correlation parameter C for the simulations as well. The frame in Figure 2.41 is perfectly ordered and has C = 1. In Figure 2.42, C ranges from 0.90 to 0.96 which is still quite high.

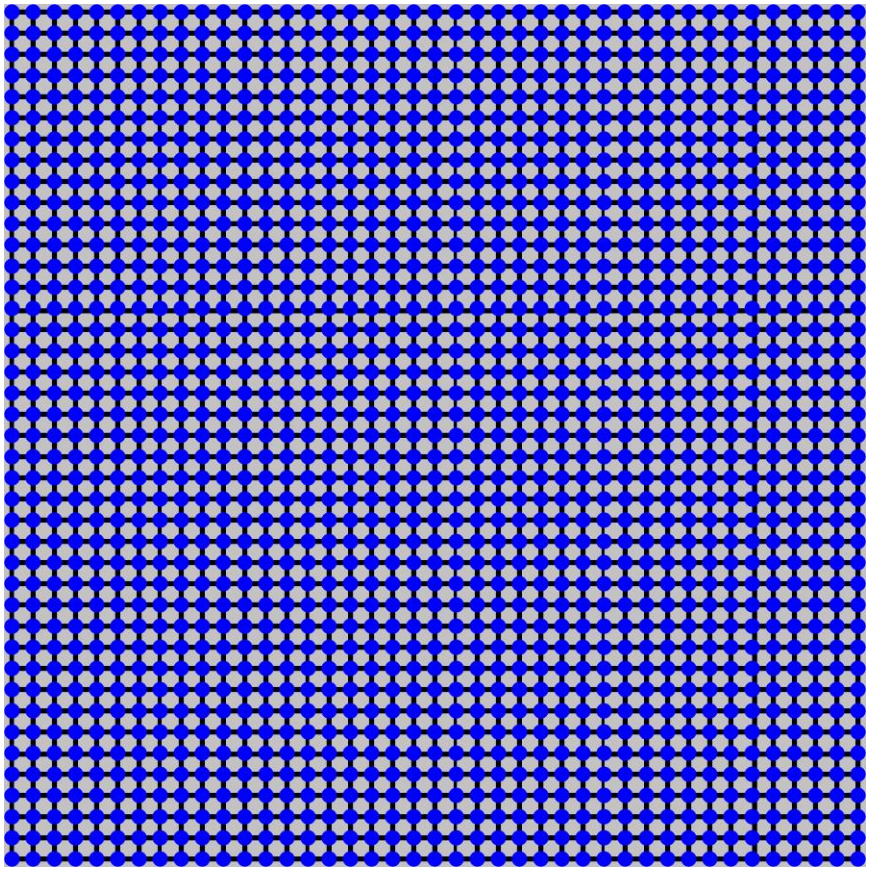

Figure 2.41 **No disorder square Monte Carlo**. With no disorder, square simulations produce perfect $C = 1$ ordering. This frame is taken after 800,000 Monte Carlo steps.

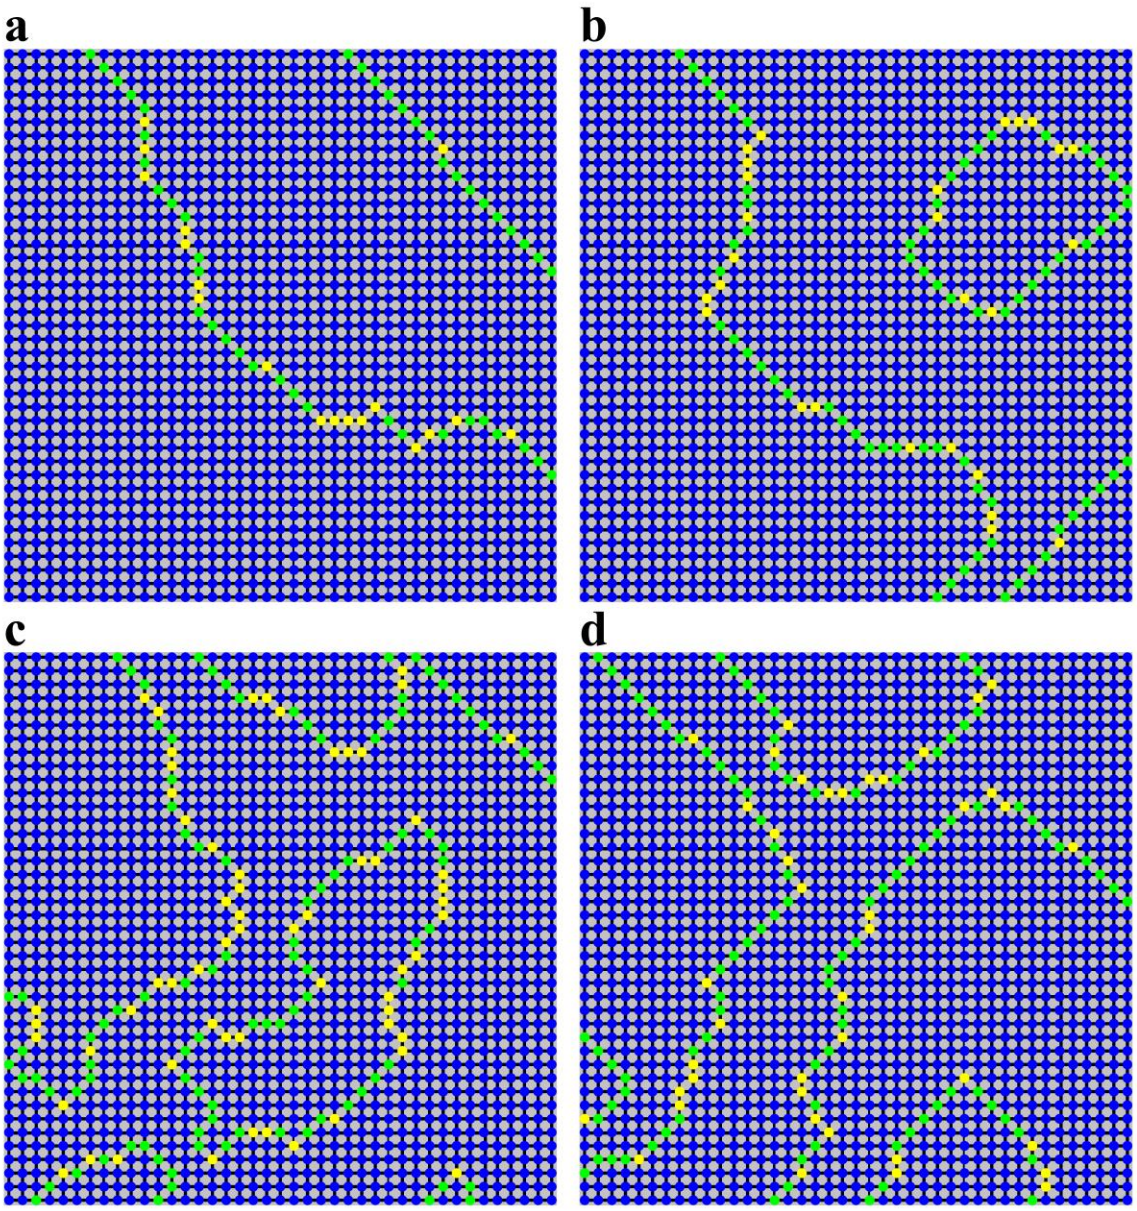

Figure 2.42 **Square Monte Carlo with varying disorder**. Final frames of square kinetic Monte Carlo simulations with domain patterns that are remarkably similar to those observed in experiment. The spread in widths used in these runs are 0.5 nm (**a**), 1 nm (**b**), 1.5 nm (**c**), and 2.5 nm (**d**). The ordering parameter *C* values are 0.96, 0.94, 0.90, and 0.92, respectively. Each image is the state of the system after 800,000 Monte Carlo steps, or an average of 200 flips per spin.

Sometimes samples can have highly disordered domains and little long-range order. This is generally due either to the elements being too small with low moments, the $FePd_3$ film having a low quality from deposition, or there being strong non-Gaussian disorder in

the system from lithography. Strong disorder, for example, can come from one element that is fabricated very wide or very thin or some other kind of non-Gaussian defect in the system. Vacancies typically do not greatly affect the ordering ability of experimental samples, but other types of quenched disorder can. Figure 2.43 shows a simulation with a higher degree of disorder and correspondingly smaller ground-state domains. This much disorder was rarely observed in real square geometry samples, but was sometimes observed in the topological defect samples discussed in Chapter 4. The *C* value for the lattice in Figure 2.43 is 0.80, which is still pretty high. As we will see, this is significantly higher than the highest observed correlation parameters for the kagome geometry. Even with what appears to be a relatively disordered system, there are still some very large regions of ground-state order.

The simulations start in a random state and due to the benefits of the kinetic Monte Carlo technique, find highly ordered states very quickly. Figure 2.44 shows snapshots of the system throughout the simulation. The lattice becomes well-ordered after only about 2000 Monte Carlo steps in Figures 2.44(a)-(d), having formed large domains of ground-state order separated by domains walls made from type II and type III vertices. Once these domains and domain walls are formed, the domains tend to shrink slightly as seen in Figures 2.44(e)-(i). This takes 30,000 Monte Carlo steps, compared to 2000 for Figures 2.44(a)-(d), showing how quickly the system can reach a relatively well ordered configuration. After the large domains have formed, the type III vertices will simply fluctuate and move around as seen in Figures 2.44(j)-(l). This behavior and the evolution of the system is in agreement with other computational studies of thermally active square ASI [119]. I often run the simulations for up to 800,000 steps to reach an average of 200

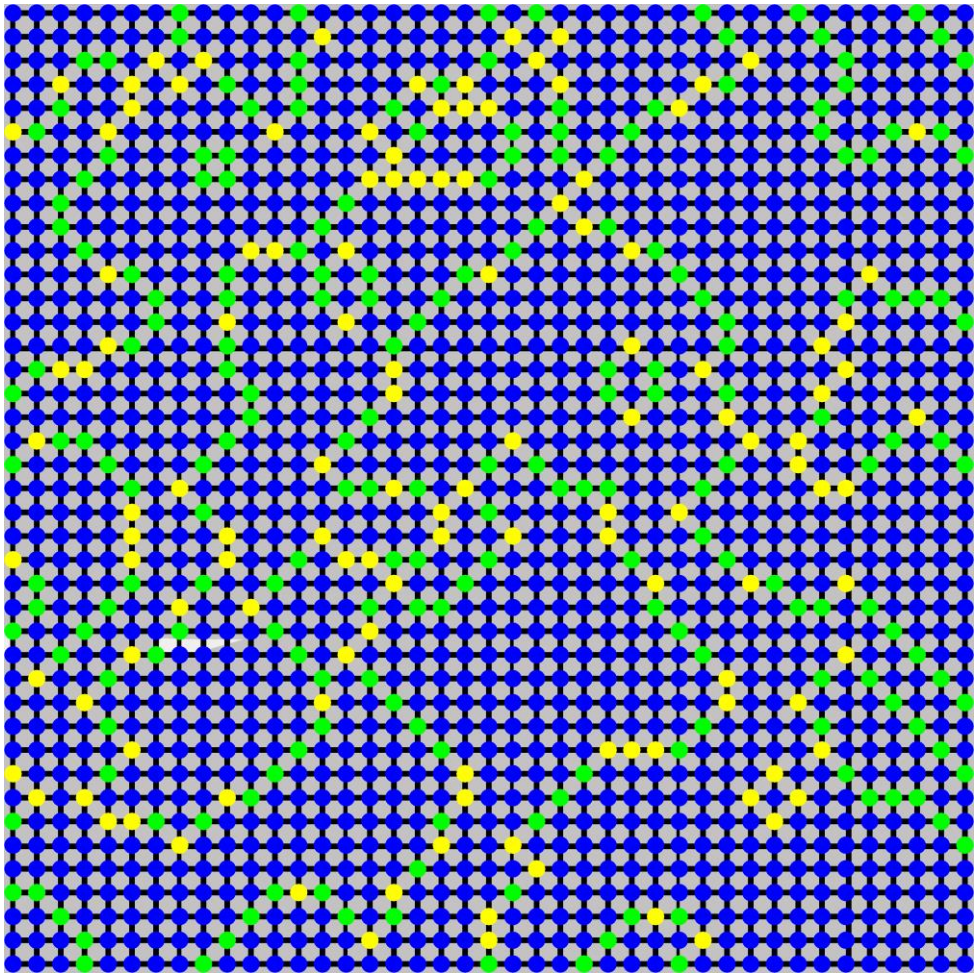

Figure 2.43 **Square Monte Carlo with high disorder**. Final frame of a kinetic Monte Carlo simulation with 3.5 nm width disorder after 800,000 Monte Carlo steps. The lattice has many smaller domains and many domain walls and *C* for this frame is 0.80.

flips per spin. Figure 2.44(l) shows the system after 40,000 Monte Carlo steps, but the behavior in Figures 2.44(j)-(l) can persist for hundreds of thousands more steps without providing much new information. Sometimes the lattice can become slightly more ordered after many steps, but generally not significantly. The *C* value starts at 0.24 in Figure 2.44(a) and increases to 0.91 in Figure 2.44(l). Overall, we are successful at modeling observed behavior in thermally active square ASI and have shown that disorder plays an important role in the long-range ordering of the system.

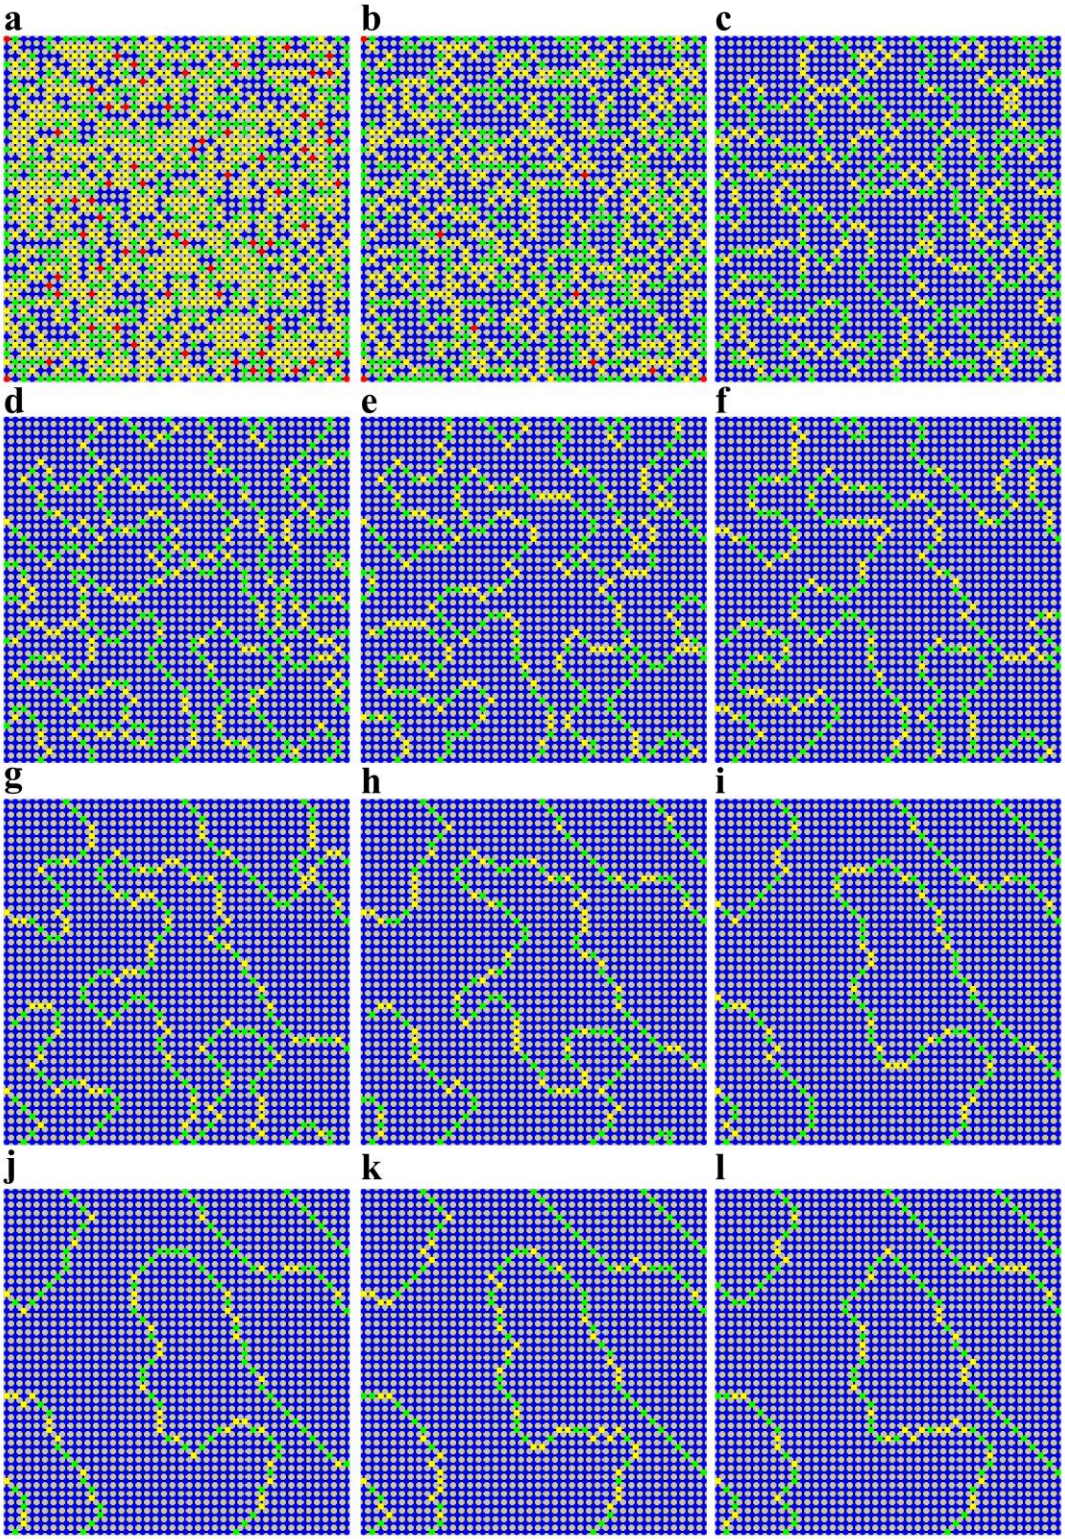

Figure 2.44 **Evolution of a square geometry kinetic Monte Carlo simulation**. Snapshots of the state of the system after (**a**) 200, (**b**) 600, (**c**) 1200, (**d**) 2000, (**e**) 4000, (**f**) 8000, (**g**) 12000, (**h**) 20000, (**i**) 30000, (**j**) 32000, (**k**) 36000, (**l**) 40000 Monte Carlo steps.

Chapter 3 Charge-ordering in artificial kagome ice

3.1 Introduction

This chapter aims to investigate theoretical predictions of long-range charge-ordering in kagome ASI and increase the body of knowledge surrounding this fully frustrated system. Unlike the square geometry, which has only a two-fold degenerate ground state, the kagome geometry is extensively frustrated with more degeneracy and a much richer predicted phase diagram. Locally, the three-island kagome vertices are six-fold degenerate, though it is predicted that long-range dipolar interactions can break this degeneracy and realize more-ordered states. As was shown in Figure 1.23, the theoretically predicted phase diagram for the kagome system has three plateaus in entropy corresponding to three distinct states of varying degeneracy: the ice rule obeying kagome I state, the charge-ordered kagome II state, and a spin-ordered state.

This chapter will focus on the realization, experimental investigation, and numerical modeling of the kagome II charge-ordered state, shown schematically in Figure 1.24(c). The chapter is divided into three main sections. Section 3.2 looks at small kagome crystals with varying edge types. Section 3.3 investigates charge-ordering in large crystals as a function of the lattice spacing and element sizes. Section 3.4 uses the kinetic Monte Carlo technique described in Section 2.5 to help better understand the behavior of the large crystals in Section 3.3. In all experimental studies, a great deal of data was collected and averaged, allowing us to approach what is likely the most natural and typical behaviors of the system.

Ice-rule ordering in the kagome geometry has been extensively observed in numerous studies since the introduction of the lattice as an ASI system [87,90,93,94,124]. Some studies have also seen ice-rule violating vertices such as those in Figure 1.26 or 1.34 in both thermal and athermal systems [78,80]. Long-range interactions and ordered states have been predicted for a long time but proved difficult to observe with rotation demagnetization protocols [93,125,127,132,191]. With the recent introduction of thermal relaxation in ASI systems, regions of charge-ordering have been observed, though the entropy of the system and the extent of the ordering is still far from the predicted states. Here, we improve on the previously reported the degrees of ordering and describe why the systems are not reaching the perfectly ordered configurations.

Perfect charge-ordering is predicted for the kagome II state, but this is not observed in practice. Instead, crystals tend to form domains of charge-ordering separated by domain walls. In Figure 3.1, I introduce a notation for describing the regions of ordering. The kagome lattice has two inequivalent lattice sites, denoted "*A*" and "*B*", as well as two-fold degenerate charge-ordering – the +*Q* charges can be on either *A* or *B* sites. If a +*Q* charge is on an *A* site or a -*Q* charge is on a *B* site, this constitutes one type of domain which I will define to be "blue" and overlay a blue color to represent it. If a -*Q* charge is on an *A* site or a +*Q* charge is on a *B* site, I will define this to be a "red" domain and overlay a red color to identify it, as shown in Figure 3.1(b). The actual colors are formally irrelevant, as the important factor is the relative sizes of the different domains. Nevertheless, I keep the colors consistent in all figures throughout the chapter. An example of my charge domains overlaid on a Lorentz TEM image is shown in Figure 3.3(c).

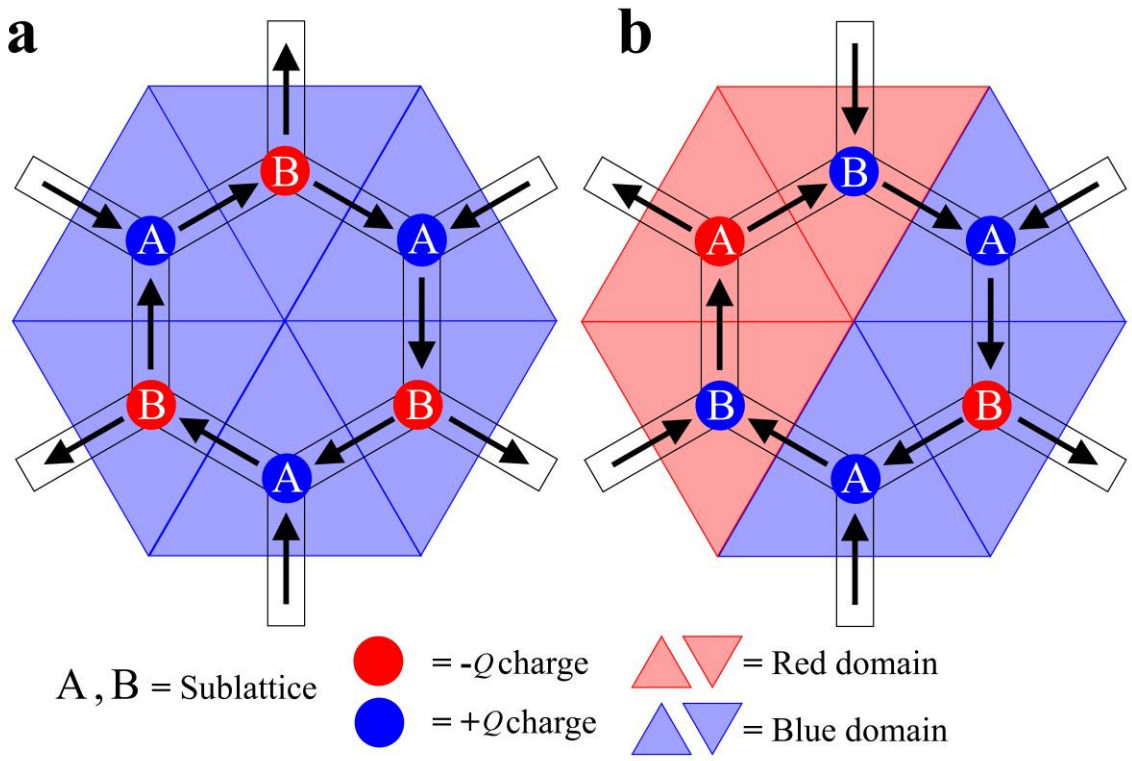

Figure 3.1 **Kagome charge domains notation**. The kagome lattice has two inequivalent lattice sites defined here as "*A*" and "*B*." If a +*Q* charge is located on an "*A*" site, I will overlay a blue domain on that vertex. Similarly, if there is a -*Q* charge on a "*B*" site I will color that with a blue domain. A +*Q* charge on a "*B*" site and a -*Q* charge on an "*A*" site will be overlaid with a red domain. **a**. Six charge-ordered vertices overlaid with a blue domain. **b**. Two domains of charge-ordering. A domain wall is created when two like charges sit on adjacent vertices. This notation is used throughout Chapter 3. In general, the red or blue domain color assigned is arbitrary and they can be interchanged without losing any information. The relative sizes of the different domains is what matters most in investigating the long-range ordering behavior of the system.

We can also quantify the degree of charge-ordering with a parameter I will call "*C*." *C* is found by assigning a 1(-1) to each pair of nearest neighbor vertices if they have the opposite (same) charge, adding up the contributions from all pairs and then dividing by the total number of nearest neighbors (see Figure 3.2). We note that $C = 1$ constitutes perfect ordering, $C = 0$ is disordered, and $C = -1$ is anti-ordering. This correlation coefficient calculation is conventional for ASI studies [75,78]. We also calculate the net magnetization

of all of our crystals. This is important because the charge-ordered state is necessarily also demagnetized, meaning the net magnetization of the sum all the elements is near zero. If a crystal is polarized along one sublattice, that crystal will also be perfectly charge-ordered, but this is trivial and is not the kagome II state we are interested in.

A polarized and trivially charge-ordered crystal is given in Figure 3.3. The magnetization of each crystal can actually be calculated along a number of directions. These include each sublattice direction, M_1, M_2, and M_3, as well as the Cartesian x and y directions, M_x and M_y. We typically use the x direction, the horizontal axis in the figures presented in this chapter. The formula is given by

$$|M_x| = \frac{\vec{M}\cdot\hat{x}}{|M_{x,max}|} \tag{3.1}$$

This projects the magnetization along the x axis. If the crystal is polarized along the x axis, $|M_x| = 1$ and if the sample is demagnetized, $|M_x| = 0$. We find that typically the $|M_y|$, $|M_1|$, $|M_2|$, and $|M_3|$ values are all approximately equal to $|M_x|$ after annealing and so we only report one of them.

The fabrication for these crystals was described in Section 2.2 and some examples of crystal sizes and geometries studied are shown in Figures 3.4 and 3.5. We fabricated crystals as small as possible, one-hexagon, up to the largest that would still fit in the TEM field-of-view with enough resolution to yield accurate and usable Lorentz contrast images, about 28 hexagons on a side or 7000 elements. We fabricated crystals with both one-fold edges (1FE) and two-fold edges (2FE). Examples of 2FE are shown in crystal 5 in Figure 3.4(a). 2FE have two elements that meet at the edge. 1FE are shown, for example, in crystal 6 in Figure 3.4(a). Here, a single element sticks out so that there is only one element on the edge. We found that for small crystals, these edge effects had a large impact,

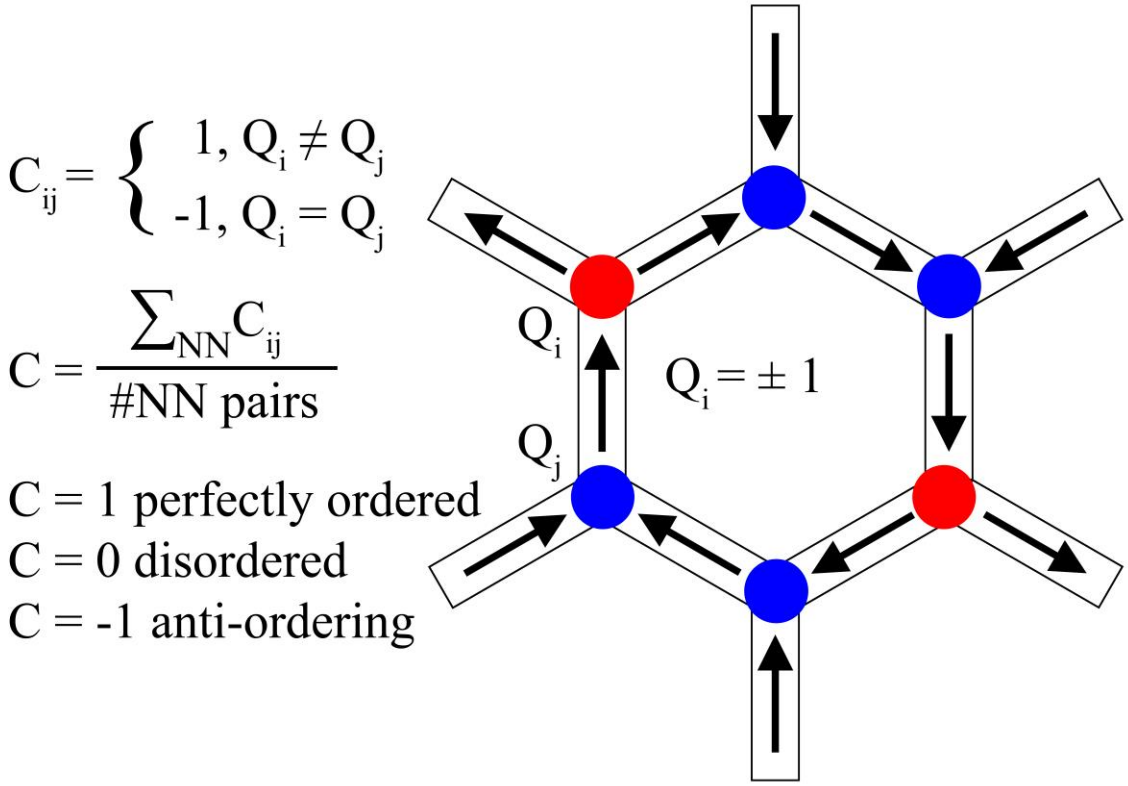

Figure 3.2 ***C* parameter calculation**. The charge correlation parameter *C* is found by assigning a 1 (-1) to a pair of neighboring vertices if they have the opposite (same) charge, adding up the contributions from all pairs and then dividing by the total number of nearest neighbors. We note that $C = 1$ constitutes perfect ordering, $C = 0$ is disordered, and $C = -1$ is anti-ordering.

but for the larger crystals, the edges did not matter. We also looked at armchair edge crystals shown, for example, in crystal 11 in Figure 3.4(b). These edges did not appreciably impact the behavior of the crystals.

To perform experiments, we first polarize the crystals as shown in Figure 3.3. This helps to check that the elements were able to flip and relax and serves as a good way to test what temperatures to heat to in the annealing protocols. A typical annealing protocol is to heat above the sample's T_C and cool at a given rate down to room temperature,

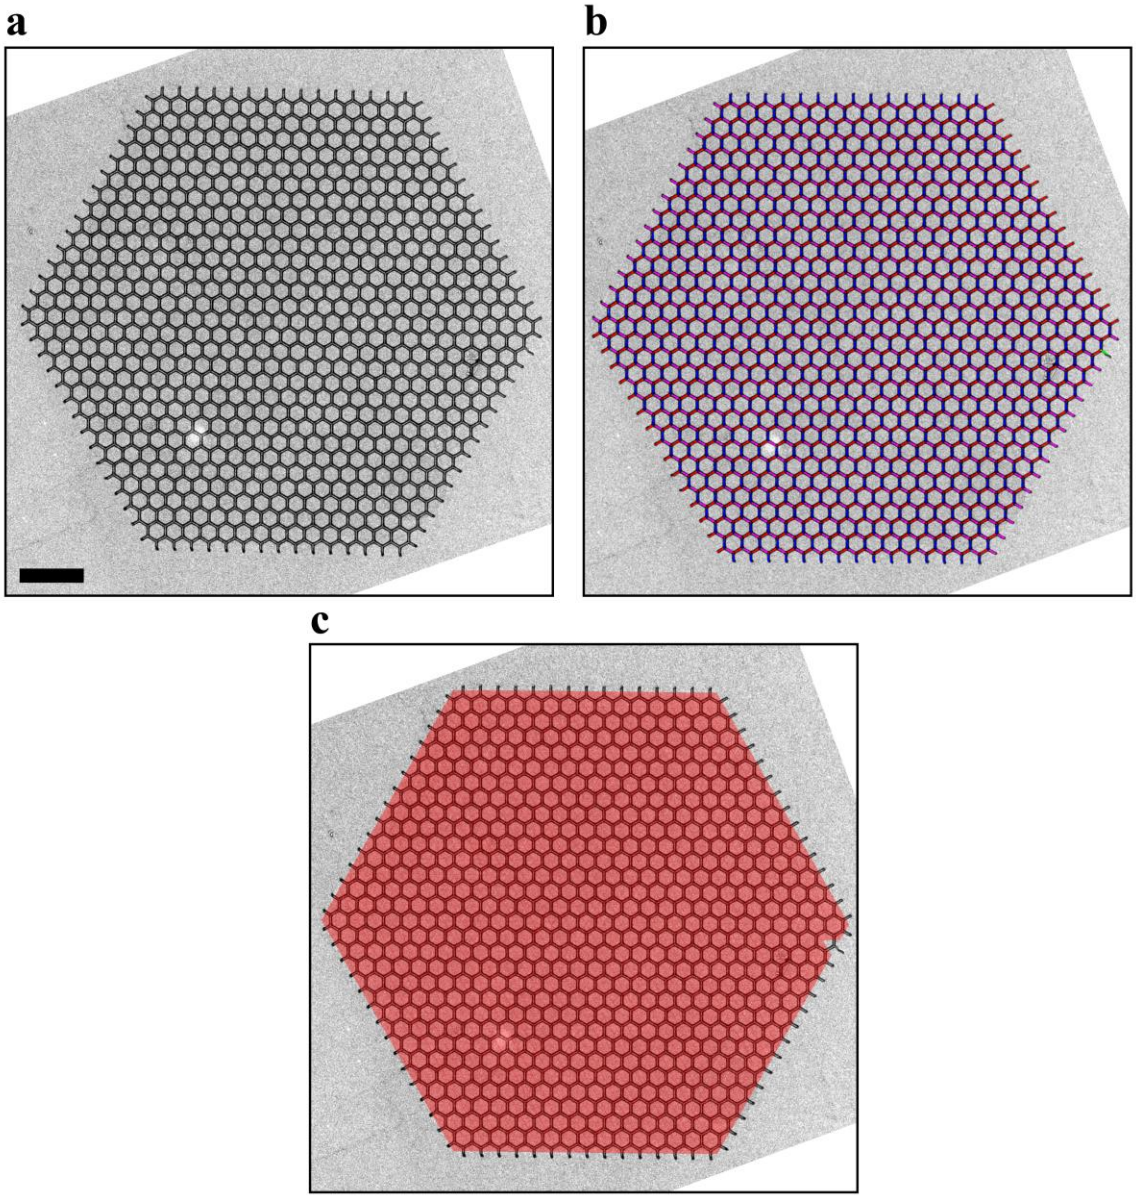

Figure 3.3 **Polarized kagome crystal. a**. Lorentz TEM image of a kagome crystal polarized toward the bottom right corner of the image. **b**. The same Lorentz TEM image as in (**a**) but with colored arrows added showing the directions of individual elements and red and blue dots added at vertices indicating the excess magnetic charge on each vertex. **c**. Charge domain map of the same crystal. The crystal is charge-ordered, but this is trivial because it is polarized. A true long-range ordered kagome ice II crystal is charge-ordered but has zero net magnetization. Charges and charge domains are defined in Figure 3.1. Scale bar is 3 μm.

as shown in Figure 2.33. For small geometry crystals, the annealing protocols vary more, so each figure with small crystals has an individual annealing graph indicating the process used for those crystals. For the large kagome crystals, the protocol is always that shown in Figure 2.33, cooling at a rate of 1 °C min^{-1}, just like the square geometry crystals already discussed in Chapter 2. After annealing, the objective lens is degaussed and Lorentz contrast images are captured like the one in Figure 3.3(a). These images are then processed and annotated, where I can add arrows on the elements and charges on the vertices, shown in Figure 3.3(b), as well as charge domain overlays, shown in Figure 3.3(b). We calculate C for each crystal as well as the net magnetization. If after annealing $|M_x|$, $|M_y|$, $|M_1|$, $|M_2|$, or $|M_3| > 0.2$, we do not use the data from that crystal as it is deemed to be too polarized and the degree of charge-ordering could be biased.

3.2 Charge-ordering and edge effects in small kagome crystals

We begin our investigation by examining the behavior and properties of the smallest possible kagome crystals after thermal relaxation and work our way up to larger and larger geometries. Understanding the behavior on these small scales aids in the understanding and theory of long-rage ordering in larger crystals. This section goes through each small kagome geometry one by one, highlighting the typically observed spin configurations and properties after annealing. Some of the first papers published on thermally active ASI systems looked at one, two, and three hexagon kagome crystals [81,84]. The small kagome geometries are interesting in their own right and yield a variety of relevant data. All element lengths in the crystals in this section are 500 nm. We will look

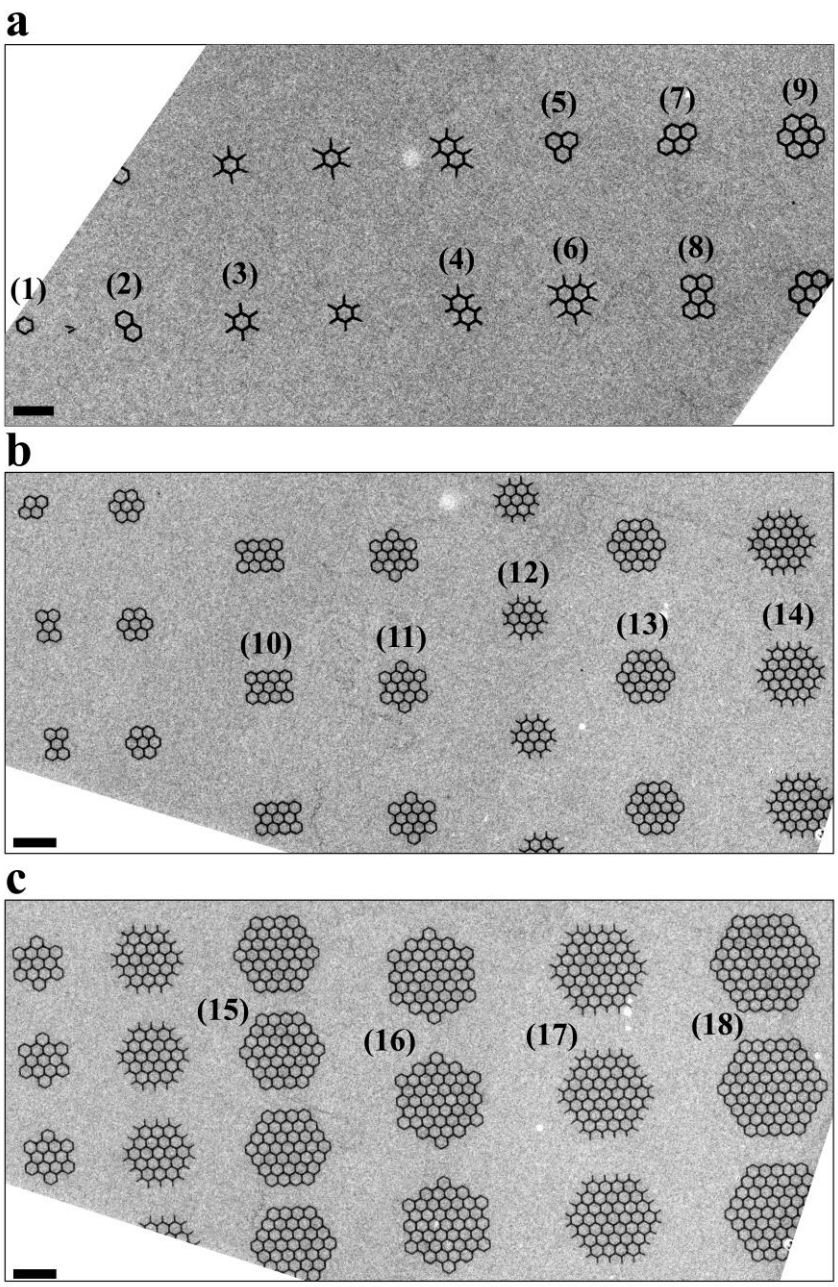

Figure 3.4 **Small kagome crystal arrays. a-c.** In-focus TEM images of various small kagome crystals studied in Section 3.2. The crystals are fabricated in arrays to take large amounts of data quickly and average results over many crystals of the same size. The crystals numbered are (1) 2FE one-hexagon, (2) 2FE two-hexagon, (3) 1FE one-hexagon, (4) 1FE two-hexagon, (5) 2FE three-hexagon, (6) 1FE three-hexagon, (7) 2FE 4-hexagon, (8) 2FE 5-hexagon, (9) 2FE 7-hexagon, (10) 2FE 11-hexagon, (11) armchair-edge 13-hexagon, (12) 1FE 7-hexagon, (13) 2FE 19-hexagon, (14) 1FE 19-hexagon, (15) 2FE 37-hexagon crystal, (16) armchair-edge 43-hexagon, (17) 1FE 37-hexagon, and (18) 2FE 61-hexagon. Scale bars for (**a**)-(**c**) are 2 μm, 3 μm, and 3 μm, respectively.

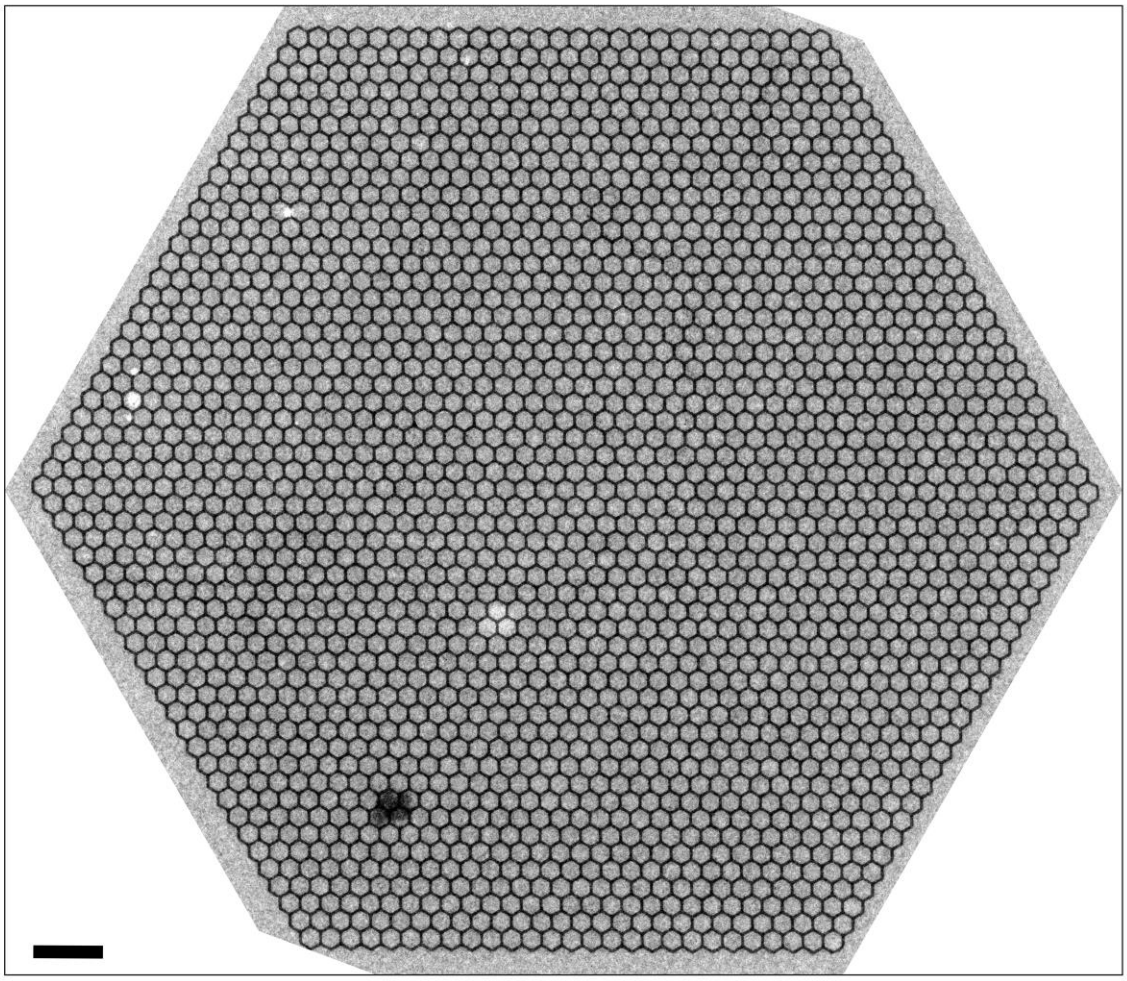

Figure 3.5 **Large kagome crystal**. In-focus TEM image of a large kagome crystal with 500 nm element lengths. Scale bar is 3 μm.

at 2FE crystals first, working our way up in size, and then compare these to 1FE crystals of similar sizes.

The smallest structure that can be studied is a single 2FE hexagon. The behavior of these crystals is shown in Figure 3.6. After annealing, these crystals are always found in either a clockwise or counterclockwise, head-to-tail, flux-closure arrangement of spins. This behavior is extremely robust and highly reproducible. Every single 2FE one-hexagon crystal we studied, approximately 100 of them, was in one of these two configurations. This is the lowest energy arrangement of spins for this geometry. Farhan *et al*. observed

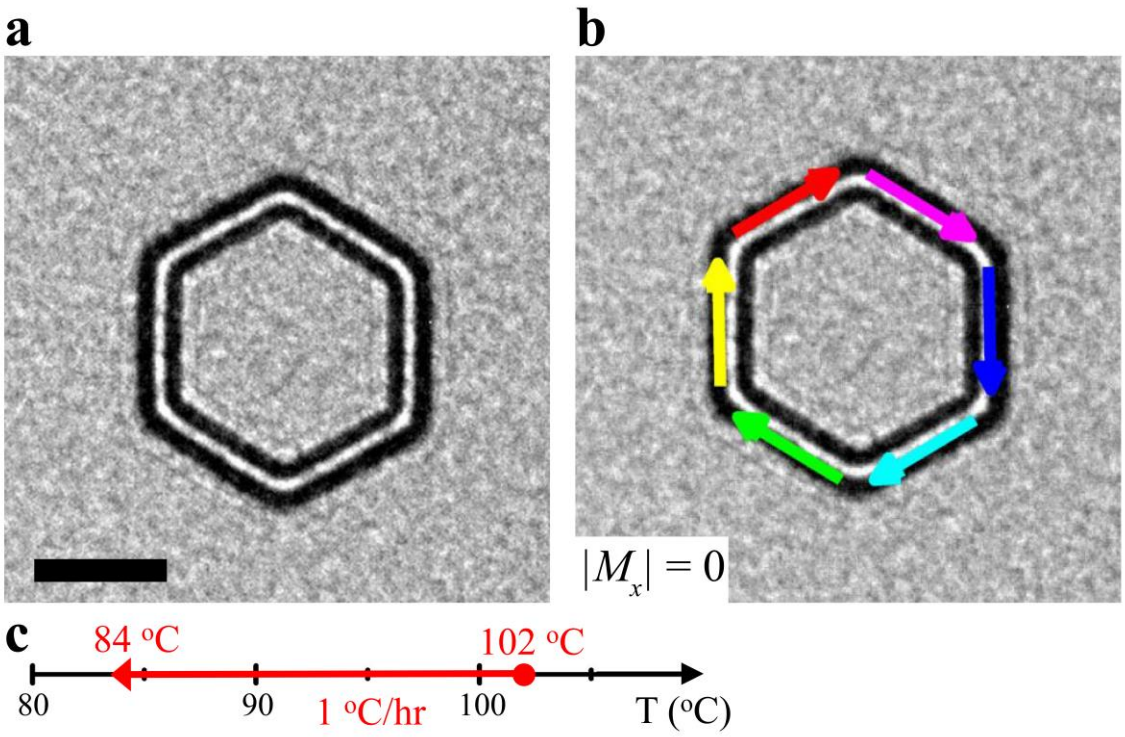

Figure 3.6 **One-hexagon 2FE crystal**. Lorentz TEM image (**a**) and corresponding spin map (**b**) of a 2FE one-hexagon kagome crystal that has been annealed. These crystals always end up in one of two states, with either a clockwise or counterclockwise head-to-tail, flux-closure arrangement of spins. $|M_x| = 0$. **c**. Annealing process for this crystal. Scale bar is 500 nm.

this state, but they also observed excitations into higher energy spin configurations with unfavorable head-to-head or tail-to-tail arrangements of spins too [81]. We note that we do not observe any other states in our samples, though we are also not in exactly the same temperature regime as Ref. [81].

Two-hexagon 2FE crystals also have extremely robust and reproducible behavior. I always found these crystals to be charge-ordered, with over 140 crystals studied. Each individual hexagon is always in a clockwise or counterclockwise flux-closure arrangement of spins. The hexagons can both have the same chirality as in Figure 3.7(c), or they can

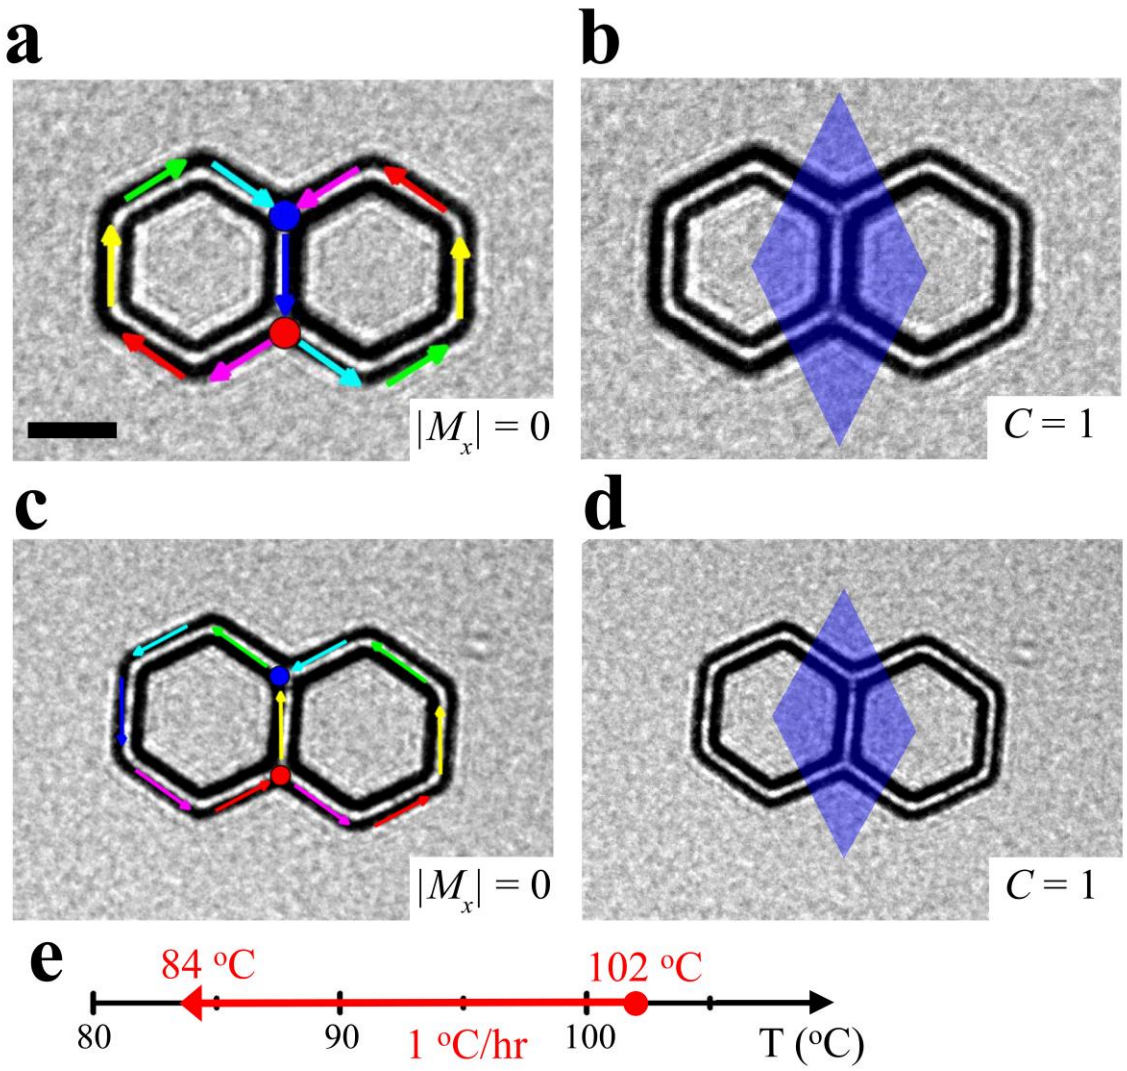

Figure 3.7 **Two-hexagon 2FE crystals. a, c.** Lorentz TEM images of annealed 2FE two-hexagon crystals with arrows and charges added. **b, d**. Lorentz TEM images with charge domain maps. **e**. Annealing process for these crystals. The 2FE two-hexagon crystals are always found to be charge-ordered. They also always are observed to have head-to-tail, flux-closure arrangements of elements on the edges. There are two qualitatively distinct states, shown here, the crystals are found in within this manifold. In (**a**), the left hexagon has a clockwise arrangement of spins and the right hexagon has a counterclockwise arrangement. In (**c**), both hexagons have a counterclockwise arrangement. Both crystals have $C = 1$ and $|M_x| = 0$. The behavior observed here is extremely robust and reproducible. Scale bar is 500 nm.

have opposite chirality as in Figure 3.7(a). In Figure 3.7(a), the middle spin fits into the head-to-tail chains of both hexagons, while in Figure 3.7(c), the middle spin only fits into the head-to-tail chain of the left hexagon. This results in the configuration found in Figure 3.7(a) having slightly lower energy than in Figure 3.7(b), but both states are still low energy, charge-ordered, and demagnetized.

In contrast, I almost never found 2FE three-hexagon crystals to be charge-ordered. They have a very strong tendency for the individual hexagons to form head-to-tail, flux-closure arrangements of spins on the edges, which is not compatible with charge-ordering in this geometry. In Figure 3.8(a), the spins on the edges all follow a counterclockwise arrangement. In order for the crystal to be charge-ordered, the three outer vertices must have a different charge from the inner one, but this is not possible for any configuration of spins that has continuous head-to-tail ordering around the edge elements of each hexagon. Since the crystal is generally not charge-ordered, it is most often found with the charge domain pattern shown in Figure 3.8(b). We have found only one three-hexagon crystal to be perfectly charge-ordered out of about 100 studied. This crystal is shown in Figure 3.9. The charge-ordering is allowed because of an energetically unfavorable head-to-head arrangement of spins in the lower left corner of the image. This is very uncommon and it is clear from this geometry that the edge effects in 2FE crystals tend to dominate other low energy or ordered configurations.

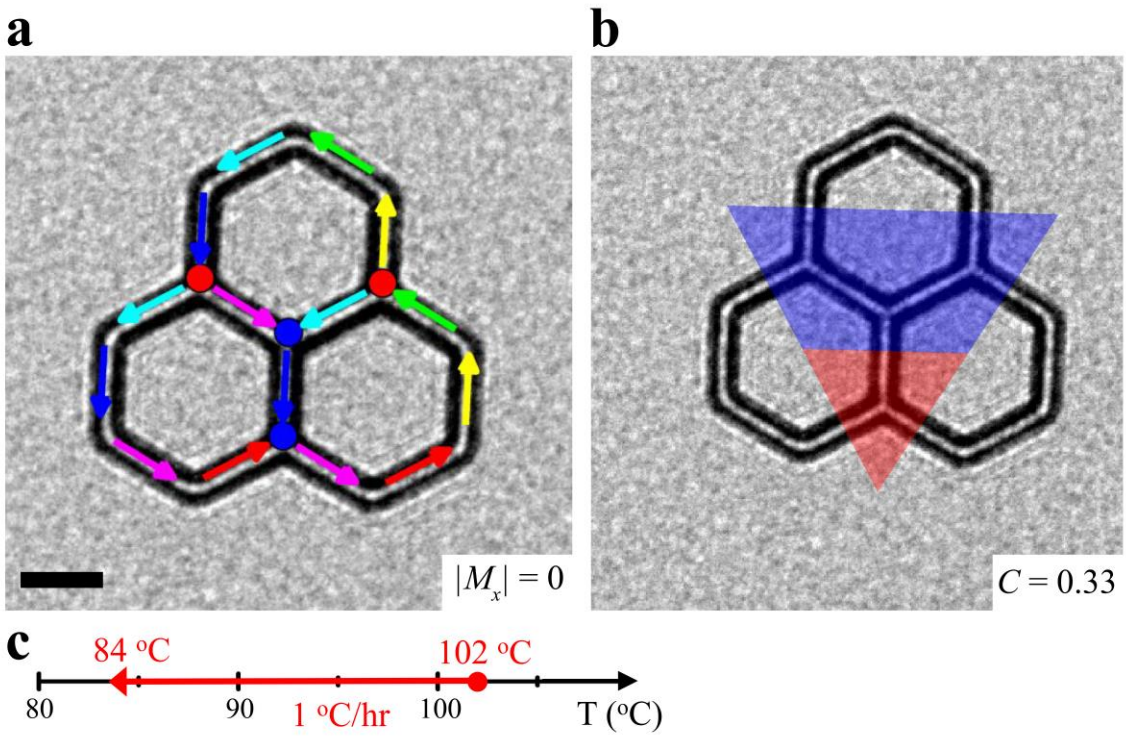

Figure 3.8 **Charge domains in a three-hexagon 2FE crystal**. **a.** Lorentz TEM image of an annealed 2FE three-hexagon crystal with arrows and charges added. **b**. Charge domain map. **c**. Annealing process for this crystal. $C = 0.33$ and $|M_x| = 0$. Three-hexagon crystals are almost never found to be charge-ordered because the charge-ordering is not compatible with having flux-closure arrangements of spins on the edges. Scale bar is 500 nm.

To solidify the energetics of the three-hexagon states we model the crystal using the Object Oriented MicroMagnetic Framework (OOMMF) package from NIST. OOMMF uses the Landau-Lifshitz-Gilbert (LLG) equation given by [192]

$$\frac{d\boldsymbol{M}}{dt} = -|\gamma|\boldsymbol{M} \times \boldsymbol{H}_{eff} + \frac{\alpha}{M_s}\left(\boldsymbol{M} \times \frac{d\boldsymbol{M}}{dt}\right) \tag{3.2}$$

where $\boldsymbol{M}$ is the local magnetization, γ is the gyromagnetic ratio, M_s is the saturation magnetization, $\boldsymbol{H}_{eff}$ is the local effective magnetic field, and α is a constant. This equation and the OOMMF package do an excellent job of treating magnetic systems as they take into account the full size and domain structures of the patterned magnets. The structure is defined in the software along with an initial configuration. The system is then allowed to

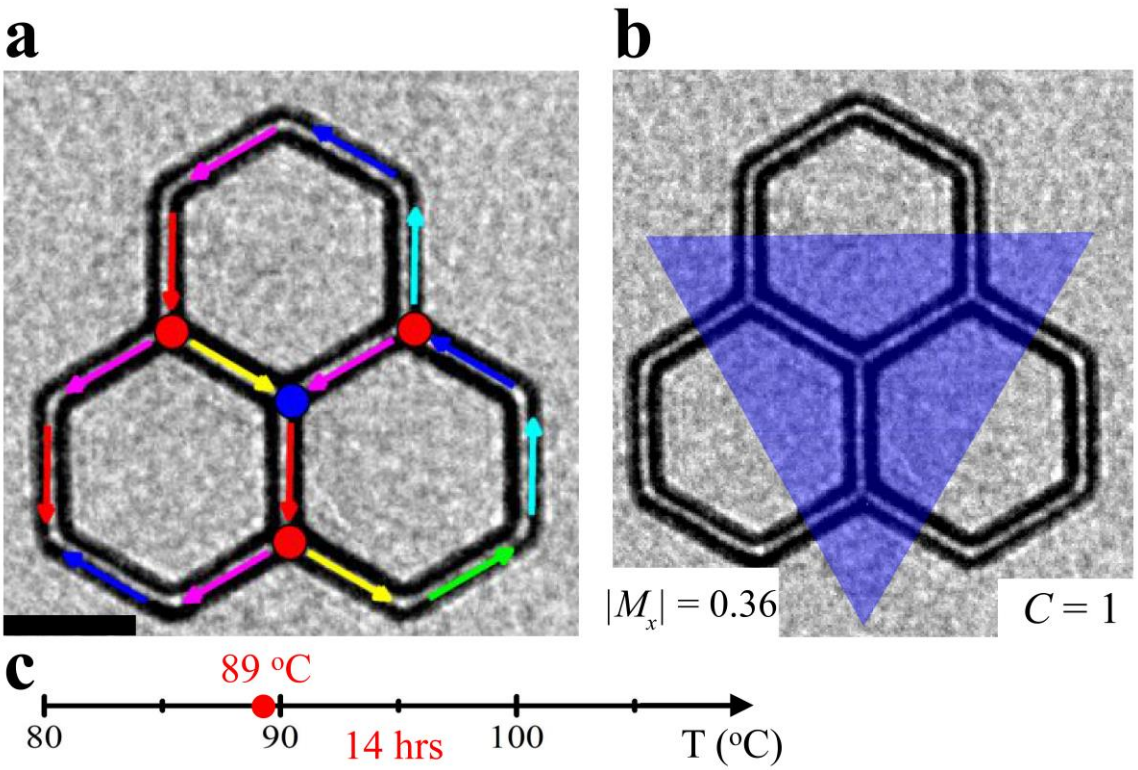

Figure 3.9 **Charge-ordered three-hexagon 2FE crystal**. **a.** Lorentz TEM image of an annealed 2FE three-hexagon crystal with arrows and charges added. **b**. Charge domain map. **c**. Annealing process for this crystal. $C = 1$ and $|M_x| = 0.36$. This is the only three-hexagon crystal we observed to be charge-ordered out of ~100. It is ordered because of an energetically unfavorable head-to-head arrangement of spins in the lower left corner. Scale bar is 500 nm.

relax stepwise through the LLG equation. Once the micromagnetic state has relaxed, the software can calculate various energies in the system. Our OOMMF results are shown in Figure 3.10. As expected from the experimental ASI results, the charge-ordered three-hexagon configuration has a significantly higher total energy than the edge flux-closure dominated state. Thus, the tendency of the three-hexagon system to minimize its energy typically results in non-charge-ordered configurations.

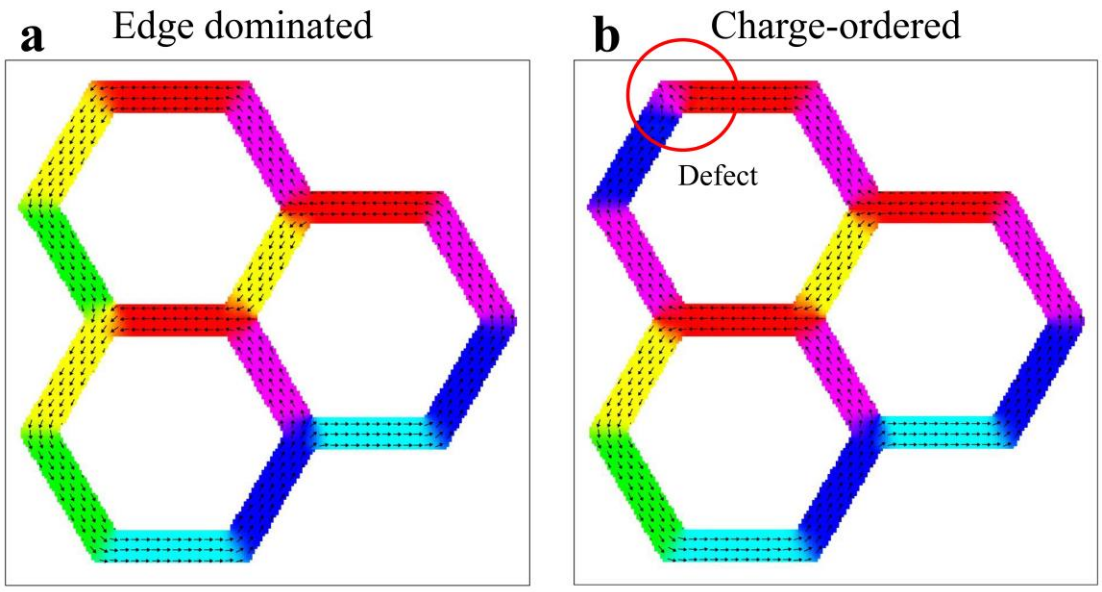

Total Energy	**2293 J/m^3**
Exchange Energy	414 J/m^3
Anisotropy Energy	100 J/m^3
Demag Energy	1779 J/m^3

Total Energy	**3043 J/m^3**
Exchange Energy	455 J/m^3
Anisotropy Energy	99 J/m^3
Demag Energy	2489 J/m^3

Figure 3.10 **Three-hexagon 2FE crystal OOMMF**. OOMMF simulation results showing the charge-ordered state (**b**) in a 2FE three-hexagon crystal has a higher total energy than the non-ordered, flux-closure state (**a**). The calculated energies for each configuration are given as well.

As the size of the crystals increase, so do the variety and complexity of their behavior. For 4-hexagon 2FE crystals, these are able to charge-order and are found to do so on some occasions, but not always. There is a spectrum of ordering and domain patterns observed for this geometry. Figure 3.11 shows a perfectly ordered 4-hexagon crystal. We have observed 10 perfectly ordered crystals out of about 50 studied, though this prevalence may be subject to some sampling bias. We did not capture images of every 2FE 4-hexagon crystal we fabricated. Still, the occurrence of both perfectly ordered and non-ordered crystals shows that the edge effects have become less important and the system tries to minimize its energy and find ordered states, though this is not always possible due to the

increasing number of possible configurations. These 4-hexagon crystals have $2^{19} = 524{,}288$ possible states. Figure 3.12 shows representative examples of other charge domain patterns observed. Figure 3.12(a) is almost perfectly anti-ordered, while Figure 3.12(b) is almost perfectly ordered. Figures 3.12(c) and (d) are moderately ordered, but illustrate the variety and diversity of possible patterns.

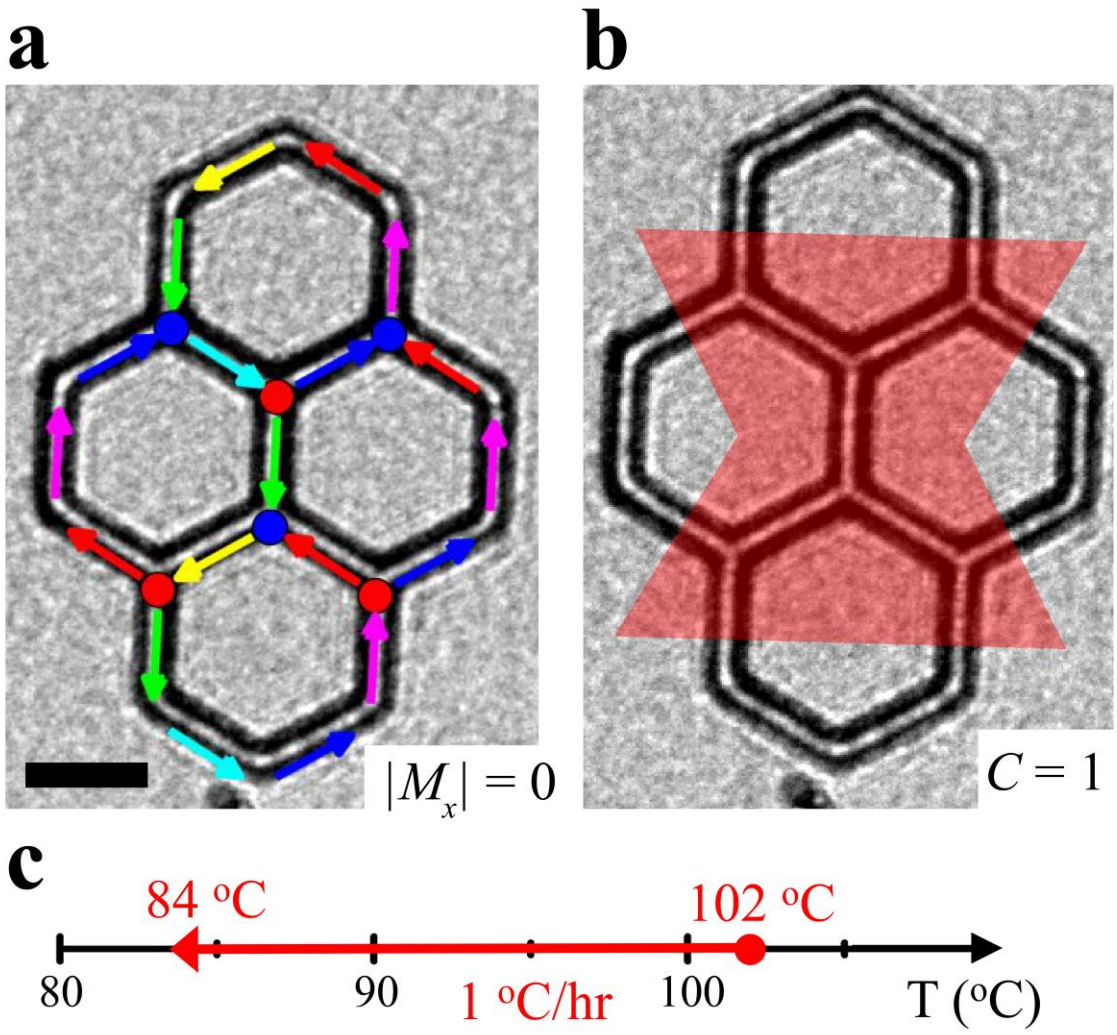

Figure 3.11 **Charge-ordered 4-hexagon 2FE crystal**. **a.** Lorentz TEM image of an annealed 2FE 4-hexagon crystal with arrows and charges added. **b**. Charge domain map of the same crystal. **c**. Annealing process for this crystal. $C = 1$ and $|M_x| = 0$. 4-hexagon crystals can order and are sometimes found to be perfectly charge-ordered, though the majority of the time they are not (see Figure 3.12). Scale bar is 500 nm.

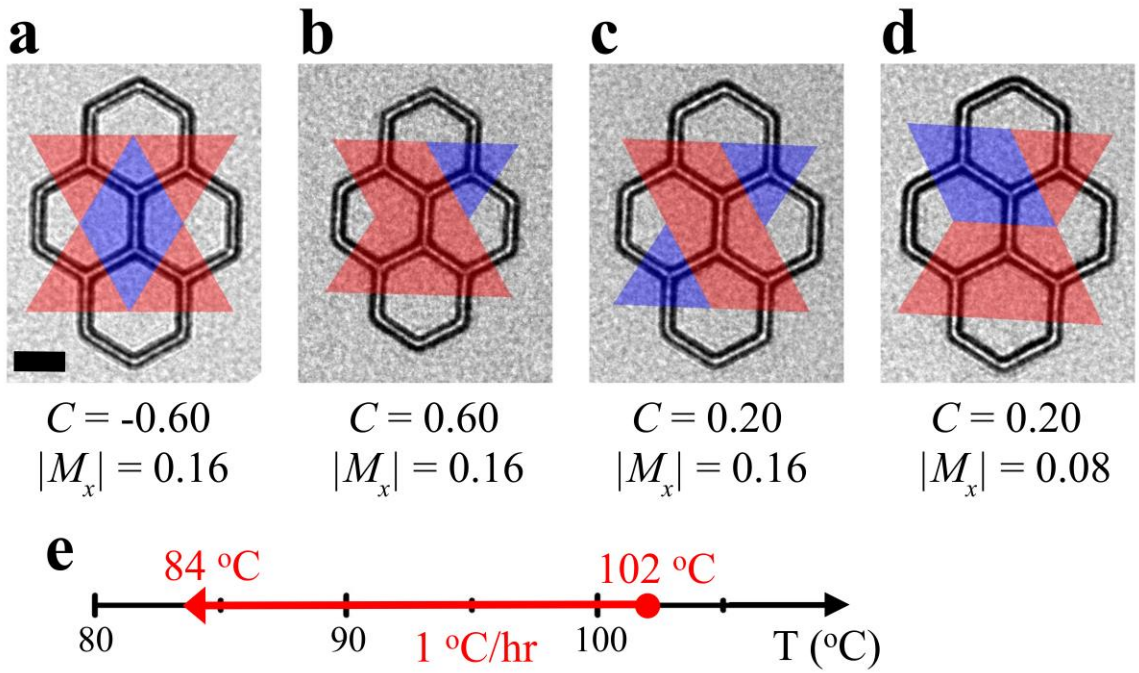

Figure 3.12 **Charge domains in 4-hexagon 2FE crystals. a-d**. Charge domain maps on Lorentz TEM images of 4-hexagon 2FE crystals. **e**. Annealing process for these crystals. These images are representative of typical observed charge domain patterns for the 4-hexagon crystals. C and $|M_x|$ are indicated for each crystal. Scale bar is 500 nm.

2FE 5-hexagon crystals are very similar to the 4-hexagon geometries. We find these crystals can also perfectly charge-order, but generally do not, due to the large number of available low-lying states with nearly the same energy. A perfectly charge-ordered crystal is shown in Figure 3.13. We have found only two perfectly charge-ordered crystals out of about 40 studied, though again there may be some sampling bias in these numbers. Some representative charge domain patterns are given in Figure 3.14. These show the diversity of the observed configurations. It is interesting to note that for Figure 3.14(a), there is only one red domain vertex, but this single opposite domain brings the average correlation C down to 0.5. The C parameters tend to be more useful when looking at larger crystals because the actual values are heavily affected by the number of hexagons in the small geometries.

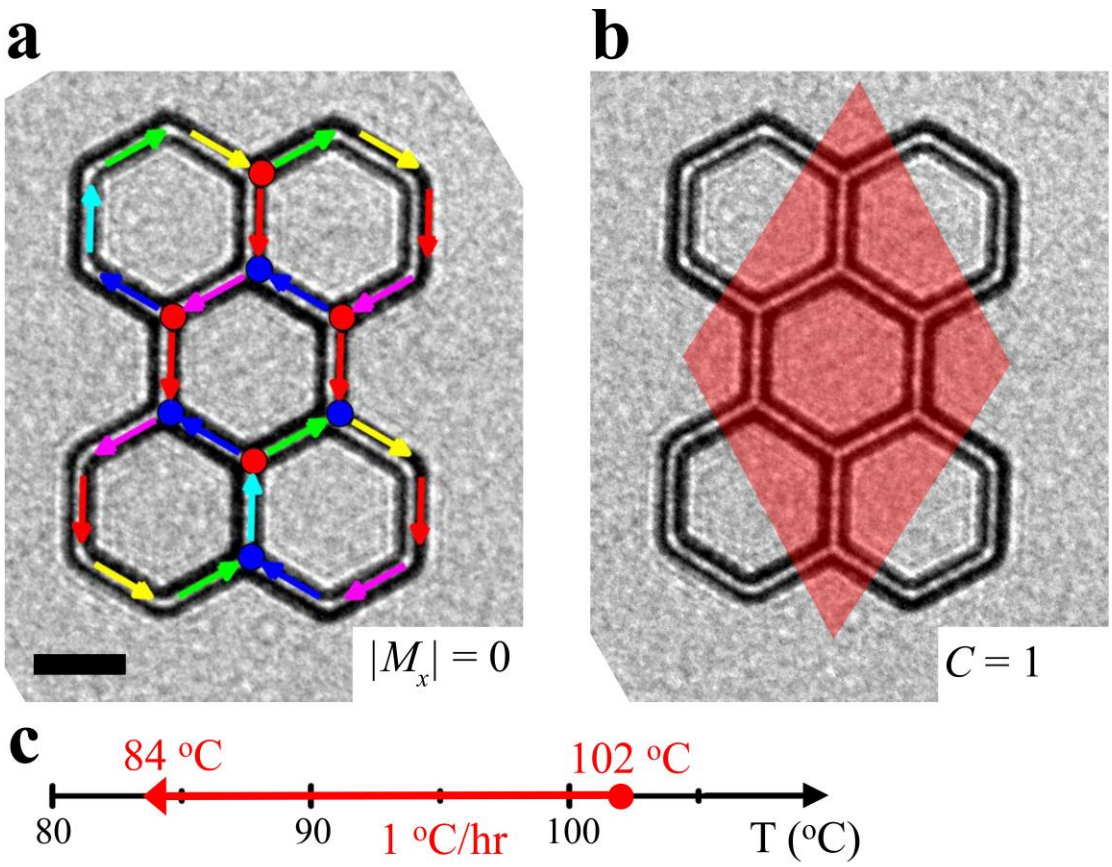

Figure 3.13 **Charge-ordered 5-hexagon 2FE crystal**. **a.** Lorentz TEM image of an annealed 2FE 5-hexagon crystal with arrows and charges added. **b**. Charge domain map of the same crystal. **c**. Annealing process for this crystal. $C = 1$ and $|M_x| = 0$. 5-hexagon crystals can order and are sometimes found to be perfectly charge-ordered, though the majority of the time they are not (see Figure 3.14). Scale bar is 500 nm.

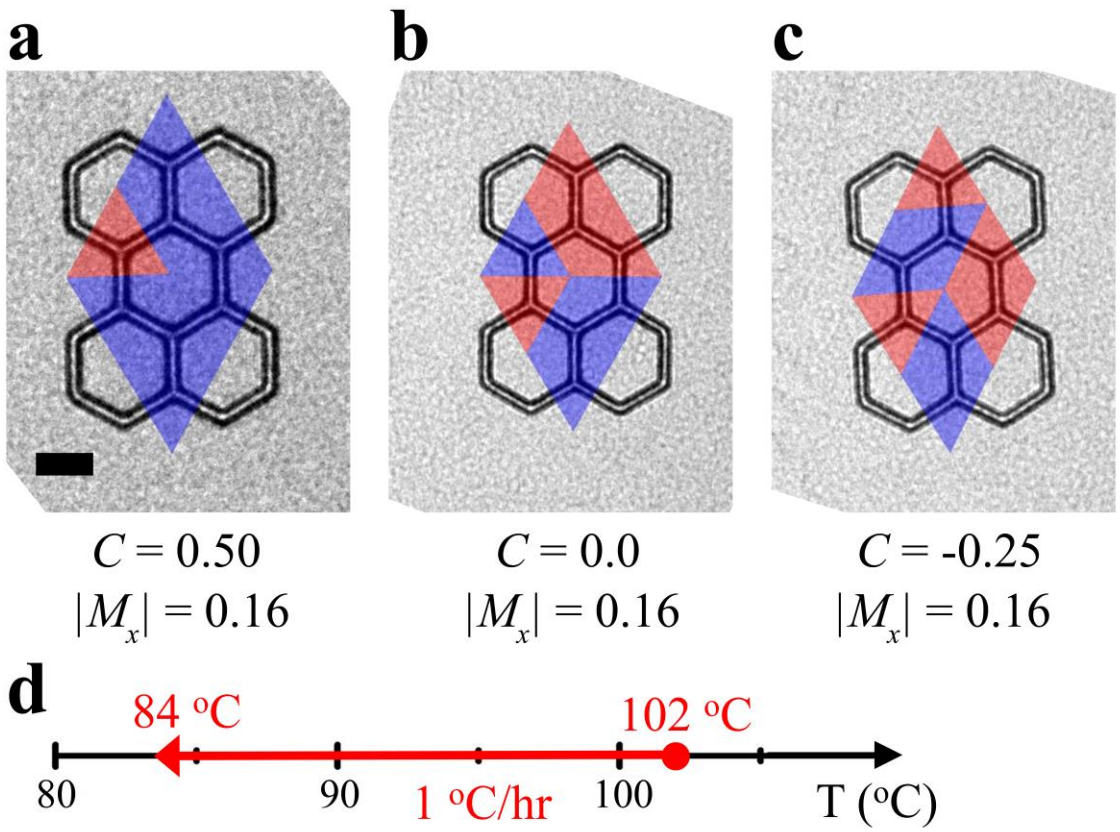

Figure 3.14 **Charge domains in 5-hexagon 2FE crystals**. **a-c**. Charge domain maps on Lorentz TEM images of 5-hexagon 2FE crystals. **d**. Annealing process for these crystals. These images are representative of typical observed charge domain patterns for the 5-hexagon crystals. C and $|M_x|$ are indicated for each crystal. Scale bar is 500 nm.

2FE 7-hexagon crystals are the largest geometry we have found to be perfectly charge-ordered. We have observed 5 of these to be perfectly ordered out of about 130 studied. A perfectly ordered 7-hexagon crystal is shown in Figure 3.15. This geometry is compatible with perfect ordering and it is interesting to note that the edges of this crystal are all still in flux-closure arrangements. Some representative examples of non-perfectly ordered 7-hexagon crystals are shown in Figure 3.16. Figure 3.16(a) has an interesting domain pattern where the central vertices belong to one domain and each of the outer vertices belong to the other. Figure 3.16(c) has a nice two-domain pattern as well.

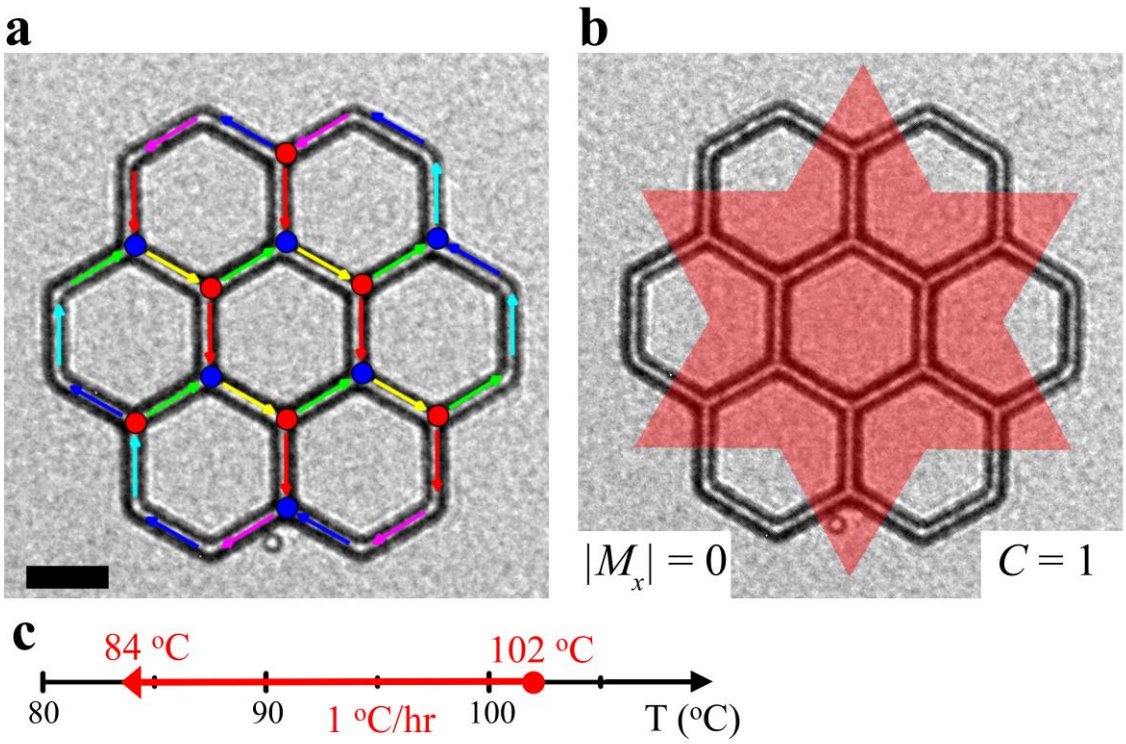

Figure 3.15 C**harge-ordered 7-hexagon 2FE crystal**. **a.** Lorentz TEM image of an annealed 2FE 7-hexagon crystal with arrows and charges added. **b**. Charge domain map of the same crystal. **c**. Annealing process for this crystal. C = 1 and $|M_x|$ = 0. 7-hexagon crystals can order and are sometimes found to be perfectly charge-ordered, though the majority of the time they are not (see Figure 3.16). 7-hexagons is the largest crystal we have observed to be perfectly ordered. Scale bar is 500 nm.

Small 1FE crystals do not have a strong desire to form flux-closure arrangements of spins on the edges. Directly comparing one-hexagon one- and two-fold edge crystals clearly shows this difference. 2FE one-hexagon crystals always find their lowest energy configurations and have extremely reproducible behavior. One-hexagon 1FE crystals, on the other hand, can be found charge-ordered as in Figure 3.17, but typically are not perfectly ordered. Figure 3.18 shows some representative domain patterns for this geometry. We have observed 27 perfectly ordered 1FE one-hexagon crystals out of about 200 studied. This is in stark contrast to the even edge crystals where 100% of them were

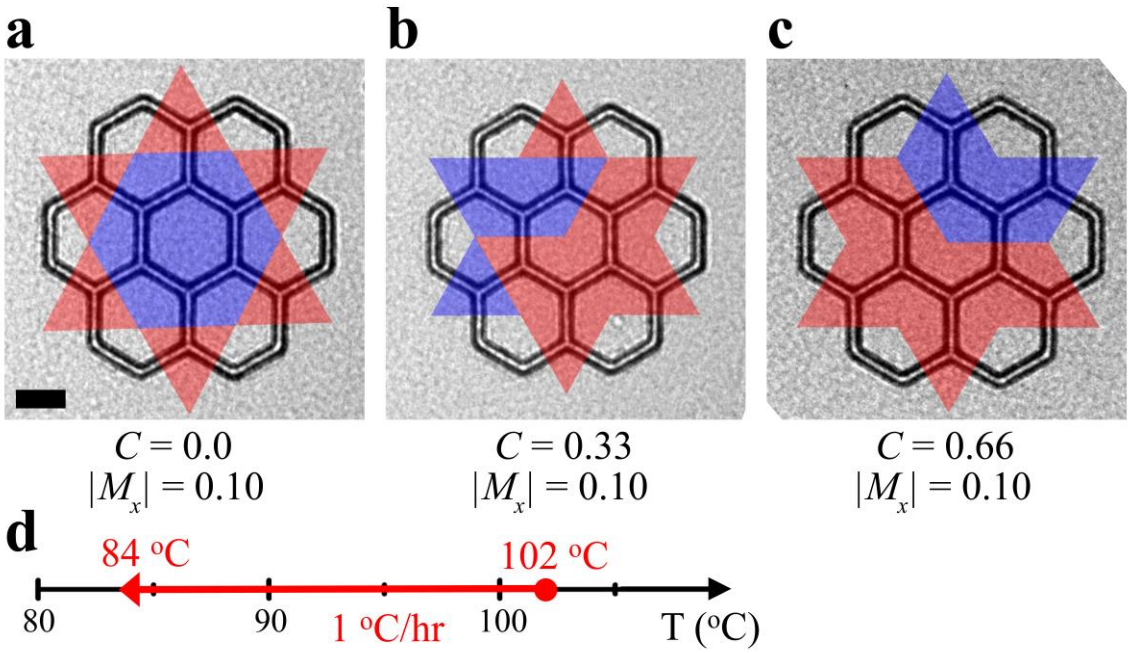

Figure 3.16 **Charge domains in 7-hexagon 2FE crystals**. **a-c**. Charge domain maps on Lorentz TEM images of 7-hexagon 2FE crystals. **d**. Annealing process for these crystals. These images are representative of typical observed charge domain patterns for the 7-hexagon crystals. C and $|M_x|$ are indicated for each crystal. Scale bar is 500 nm.

ordered. It is interesting that in Figure 3.17(a), the elements making up the inner hexagon are in a flux-closure arrangement. This is not always the case, however, as often times there can be head-to-head arrangements due to the ice-rule degeneracy of two-in/one-out or two-out/one-in three-element kagome vertices. When there are two elements meeting at an edge a head-to-head arrangement is highly unfavorable, but when there are three elements there is no energetic preference for specific configurations between the 6 degenerate vertex types.

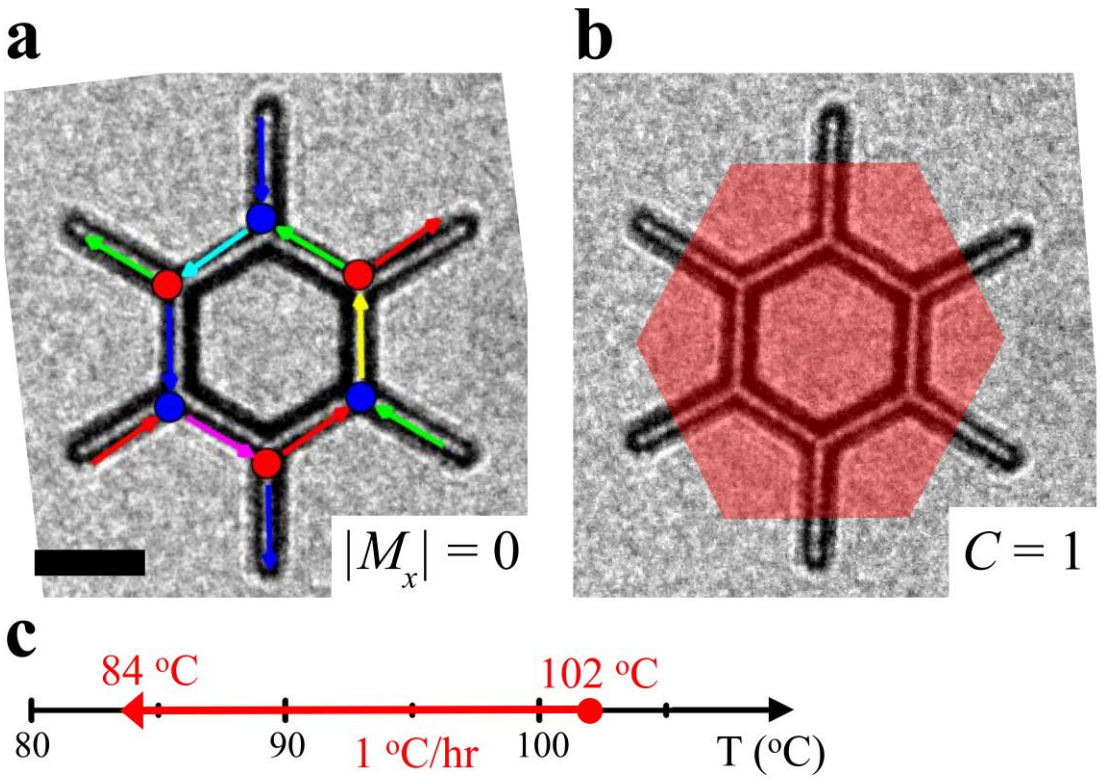

Figure 3.17 **Charge-ordered one-hexagon 1FE crystal**. **a.** Lorentz TEM image of an annealed 1FE one-hexagon crystal with arrows and charges added. **b**. Charge domain map of the same crystal. **c**. Annealing process for this crystal. $C = 1$ and $|M_x| = 0$. 1FE one-hexagon crystals can order and are sometimes found to be perfectly charge-ordered, though the majority of the time they are not (see Figure 3.18). Scale bar is 500 nm.

While 2FE two-hexagon crystals are always charge-ordered and have extremely reproducible behavior, 1FE two-hexagon crystals have much more diverse results. They are able to perfectly charge-order as is shown in Figure 3.19, but typically they are not perfectly ordered. We have only observed two perfectly charge-ordered 1FE two-hexagon crystals out of about 80 studied. Figure 3.20 shows some representative examples of charge domain patterns for these crystals. Similar to the one-hexagon case, adding 1FE breaks the flux-closure arrangements on the edges and allows for many more configurations. Adding

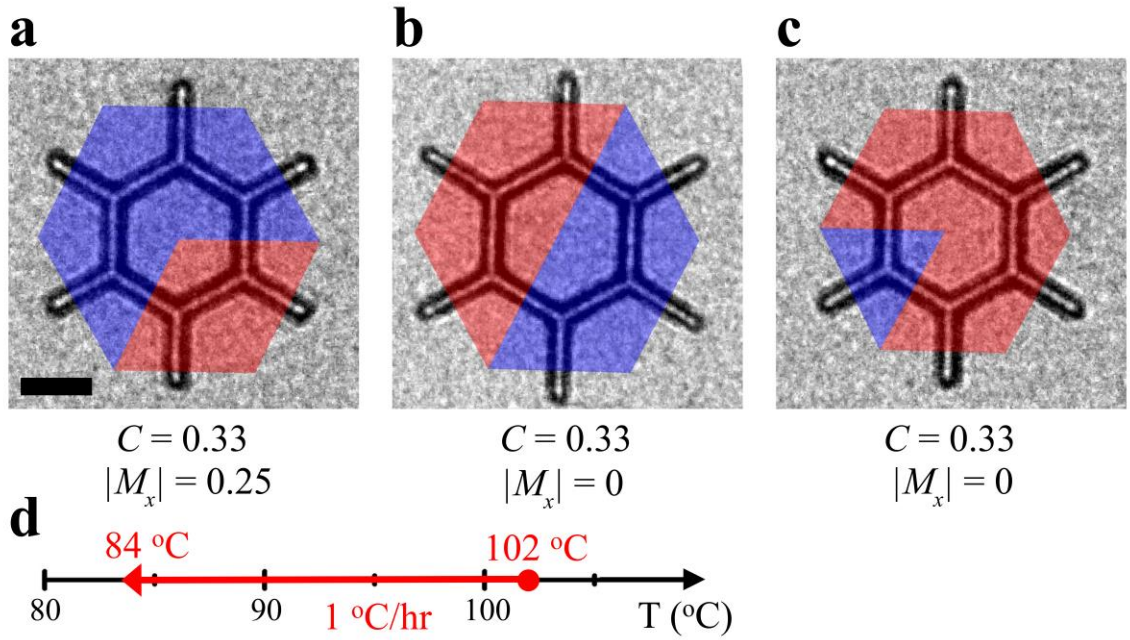

Figure 3.18 **Charge domains in one-hexagon 1FE crystals**. **a-c**. Charge domain maps on Lorentz TEM images of one-hexagon 1FE crystals. **d**. Annealing process for these crystals. These images are representative of typical observed charge domain patterns for the odd one-hexagon crystals. C and $|M_x|$ are indicated for each crystal. Scale bar is 500 nm.

the 1FE elements also simply increases the number of elements and thus the total number of possible states, making it more likely for the system to be left in a non-ordered state.

For larger crystals it is not surprising that there are many different and complex domain patterns observed. It is clear, however, that as the crystals get larger, true domains and well-ordered regions do emerge. Figure 3.21 shows some examples of 1FE three-hexagon crystals and Figure 3.22 shows a variety of assorted larger crystals that are relatively well-ordered. None of these are perfectly ordered, but some are quite close or have large domains of ordering. We note that as the crystals get larger, the edges tend not to play as much of a role in the domain patterns. This can be seen, for example, in Figures 3.22(c) and (d) or Figures 3.22(e), (f), and (g). These groups of medium sized crystals have different edges, but very similar C values and similarly sized domains. Figure 3.22(a) has a nice 1FE 7-hexagon crystal with two distinct domains. Figure 3.22(b) is an 11-hexagon

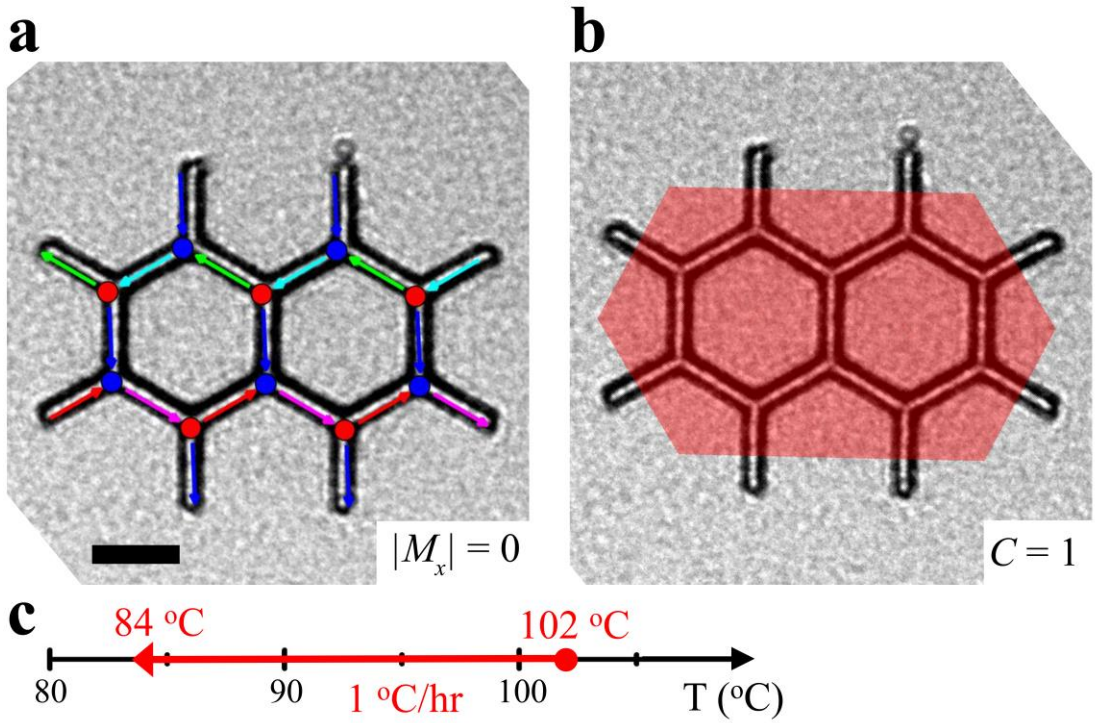

Figure 3.19 **Charge-ordered two-hexagon 1FE crystal**. **a.** Lorentz TEM image of an annealed 1FE two-hexagon crystal with arrows and charges added. **b**. Charge domain map of the same crystal. **c**. Annealing process for this crystal. C = 1 and $|M_x|$ = 0. 1FE two-hexagon crystals can order and are sometimes found to be perfectly charge-ordered, though the majority of the time they are not (see Figure 3.20). Scale bar is 500 nm.

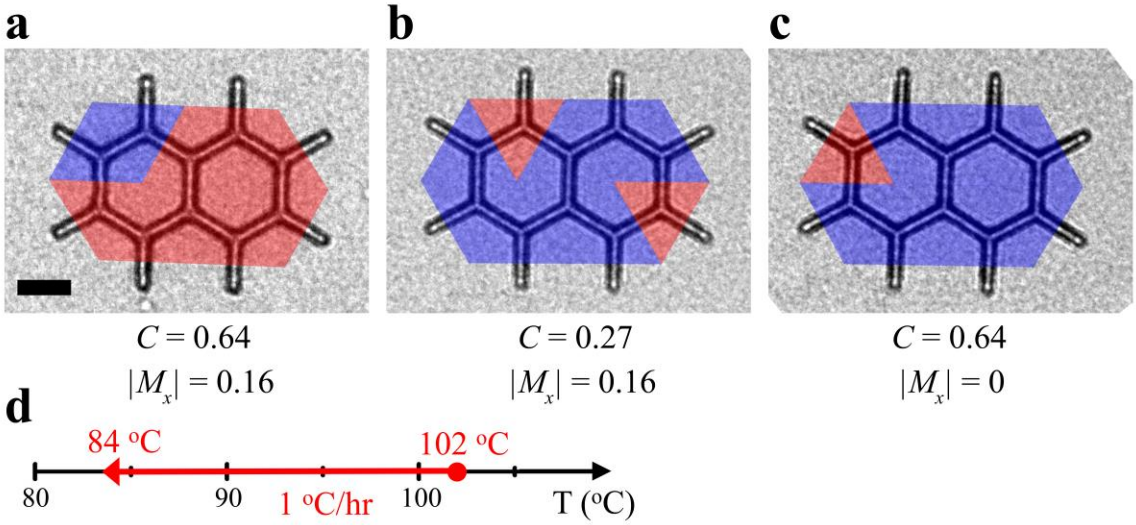

Figure 3.20 **Charge domains in two-hexagon 1FE crystals**. **a-c**. Charge domain maps on Lorentz TEM images of two-hexagon 1FE crystals. **d**. Annealing process for these crystals. These images are representative of typical observed charge domain patterns for the 1FE two-hexagon crystals. C and $|M_x|$ are indicated for each crystal. Scale bar is 500 nm.

crystal that is nearly perfectly ordered. Figure 3.22(h) has a very large blue domain that spans most of the crystal. Overall, these results and the results presented in this section show that there are a wide variety of behaviors for smaller kagome geometry crystals. Sometimes these crystals can be perfectly charge-ordered, but as their size increases, this becomes less likely and domains of ordering begin to emerge. Section 3.3 explores the charge-ordering and charge domain structure of large kagome crystals.

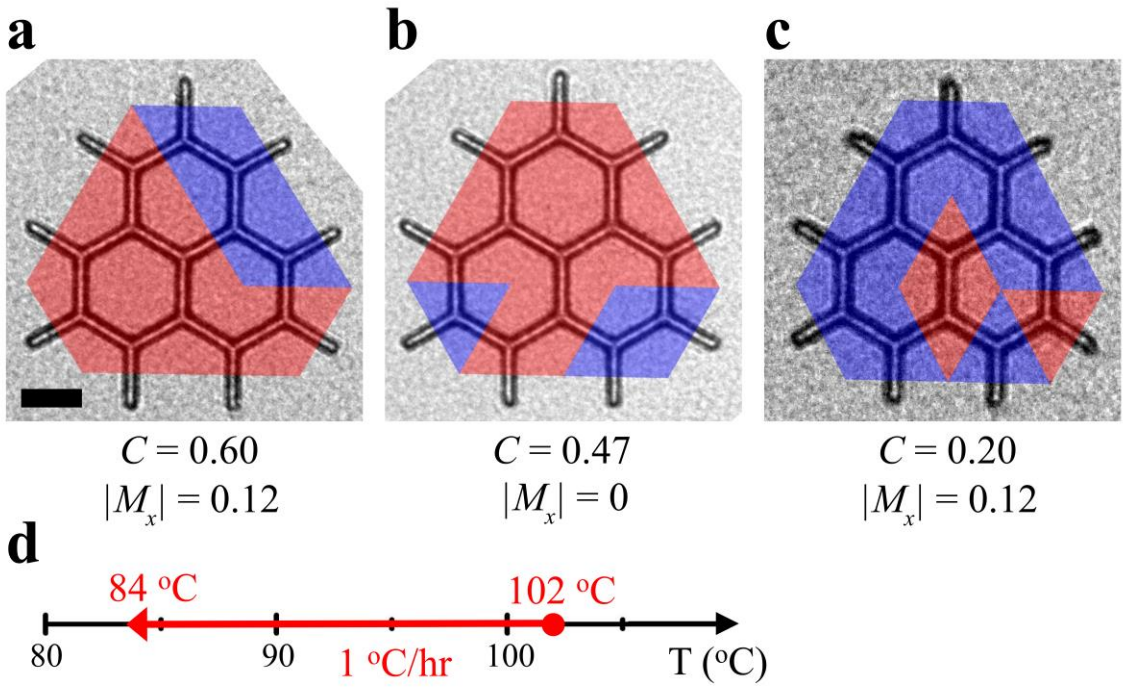

Figure 3.21 **Charge domains in three-hexagon 1FE crystals**. **a-c**. Charge domain maps on Lorentz TEM images of three-hexagon 1FE crystals. **d**. Annealing process for these crystals. These images are representative of typical observed charge domain patterns for the odd three-hexagon crystals. No perfectly charge-ordered 1FE three hexagon crystals were observed. C and $|M_x|$ are indicated for each crystal. Scale bar is 500 nm.

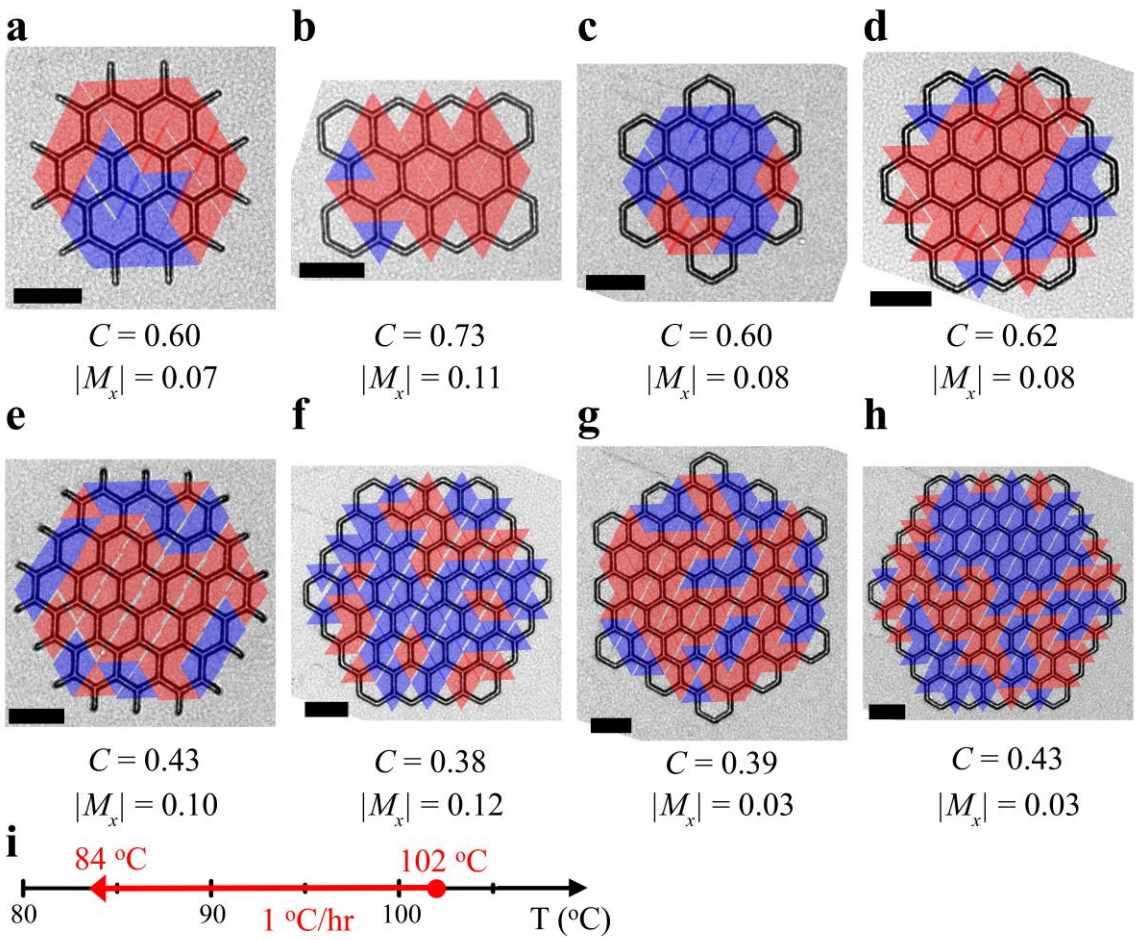

Figure 3.22 **Charge domains in assorted larger crystals**. Charge domain maps on Lorentz TEM images of varying size crystals. The crystal sizes shown are (**a**) 1FE 7-hexagon, (**b**) 2FE 11-hexagon, (**c**) armchair-edge 13-hexagon, (**d**) 2FE 19-hexagon, (**e**) 1FE 19-hexagon, (**f**) 2FE 37-hexagon crystal, (**g**) armchair-edge 43-hexagon, and (**h**) 2FE 61-hexagon. **i**. Annealing process for these crystals. These images represent a relatively well-ordered selection of larger crystals. For each size shown, other crystals of the same size were generally observed to be similarly ordered or less ordered. The crystals are all demagnetized. C and $|M_x|$ are indicated for each crystal. Scale bars are all 1 μm.

3.3 Long-range ordering in large kagome crystals

In this section we investigate the charge domain and average charge-ordering behavior of large kagome crystals. Under the same annealing conditions, the kagome geometry samples do not show the same transition to a long-range ordered state as the

square geometry. Instead, we see domains of charge-ordering similar to Ref. [78], albeit with a greater degree of ordering, which I attribute to the more relaxed temperature requirements and other benefits of $FePd_3$. This is an emergence of the predicted kagome II state [127]. In this section, as well as the next section, we examine the degree of charge-ordering as a function of element length and work to understand the factors that may be preventing the system from reaching perfectly ordered states.

The first large kagome crystals I fabricated have element lengths of 500 nm. This size was chosen mainly because it is the same size that Stephen Daunhaimer used for his ASI research in our group and we had experience fabricating samples of this size as well as well-characterized electron doses for this length. After fabrication, the samples were heated in a variety of annealing protocols to determine where elements began to flip. Most protocols yield similar results and the data presented in this section all derive from the process shown in Figure 2.33. An example of a charge domain map for a 500 nm element length crystal is shown in Figure 3.23. These 500 nm length crystals show some small domains of ordering, but they are still very far from the theoretically predicted perfectly ordered state. From observations of the behavior of the 500 nm element length crystals, we can hypothesize that using shorter elements lengths will result in higher degrees of ordering due to the increased strength of the magnetostatic interactions between neighboring spins and the magnetic charges at each vertex. To test this hypothesis, we fabricated several samples with element lengths ranging from 500 nm down to 300 nm. All of the data in this section derives from the annealing protocol in Figure 2.33, and we perform the experiments and analysis using the methods discussed in Chapter 2.

Charge domain maps of kagome samples at different element lengths are shown in Figures 3.23-3.29. Figures 3.23-3.27 show typical degrees of charge-ordering for each element length and representative domain patterns. For the shorter element lengths, some of the domains are quite large and the samples show very nice long-rage order. For the longer element lengths — the domains tend to be smaller, more disordered, have more jagged patterns, and have more point defects or few-vertex domains. There is one example each of domain patterns shown for the 500 nm, 450 nm, and 400 nm geometries. There are two examples of each of the 350 nm and 300 nm element length crystals. It is apparent from these figures that shorter element lengths tend to result in larger domains of ordering.

The best ordering we have observed is around $C = 0.6$ in the 350 nm length samples, significantly higher than was previously reported in Ref. [78]. We also note that we do not observe any energetically unfavorable ± 3 charges in our samples in contrast to other reports [78,79]. The data presented in Figures 3.26, 3.27, and 3.28 demonstrate the highest yet degree of ordering reported for the kagome geometry. Figure 3.30 also shows a spin map of the same crystal shown in Figure 3.26. This is an example of a typical spin map corresponding to any of the charge domain maps. The spins appear random and disordered, corresponding to a well-demagnetized state.

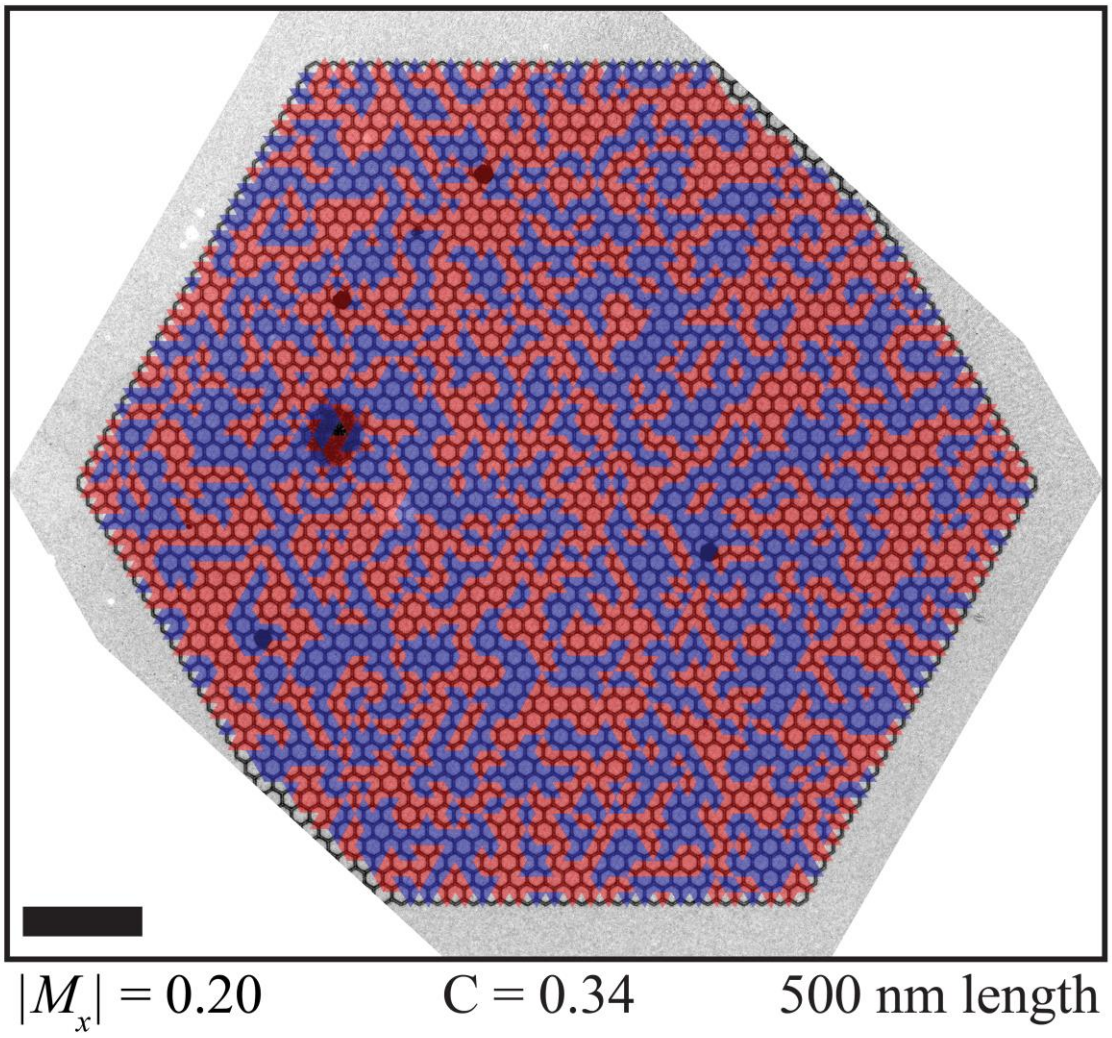

Figure 3.23 **Charge domain map of a 500 nm element length kagome crystal**. Lorentz TEM image with charge domains overlaid showing a characteristic domain pattern for 500 nm element length crystals. This crystal has $C = 0.34$ and $|M_x| = 0.20$. The average C for all 500 nm element length crystals studied is given in Figure 3.31. Scale bar is 5 μm.

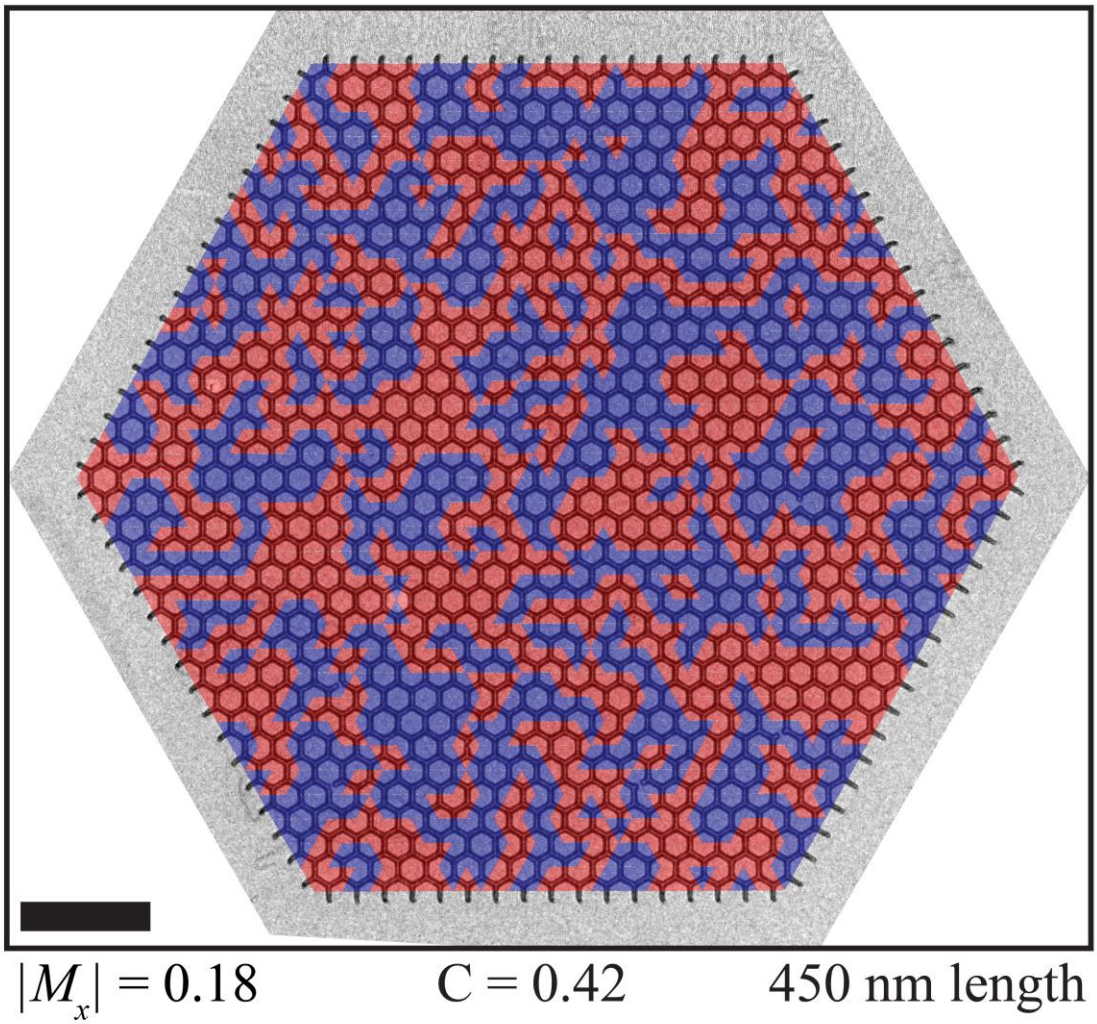

Figure 3.24 **Charge domain map of a 450 nm element length kagome crystal**. Lorentz TEM image with charge domains overlaid showing a characteristic domain pattern for 450 nm element length crystals. This crystal has $C = 0.42$ and $|M_x| = 0.18$. The average C for all 450 nm element length crystals studied is given in Figure 3.31. Scale bar is 3 μm.

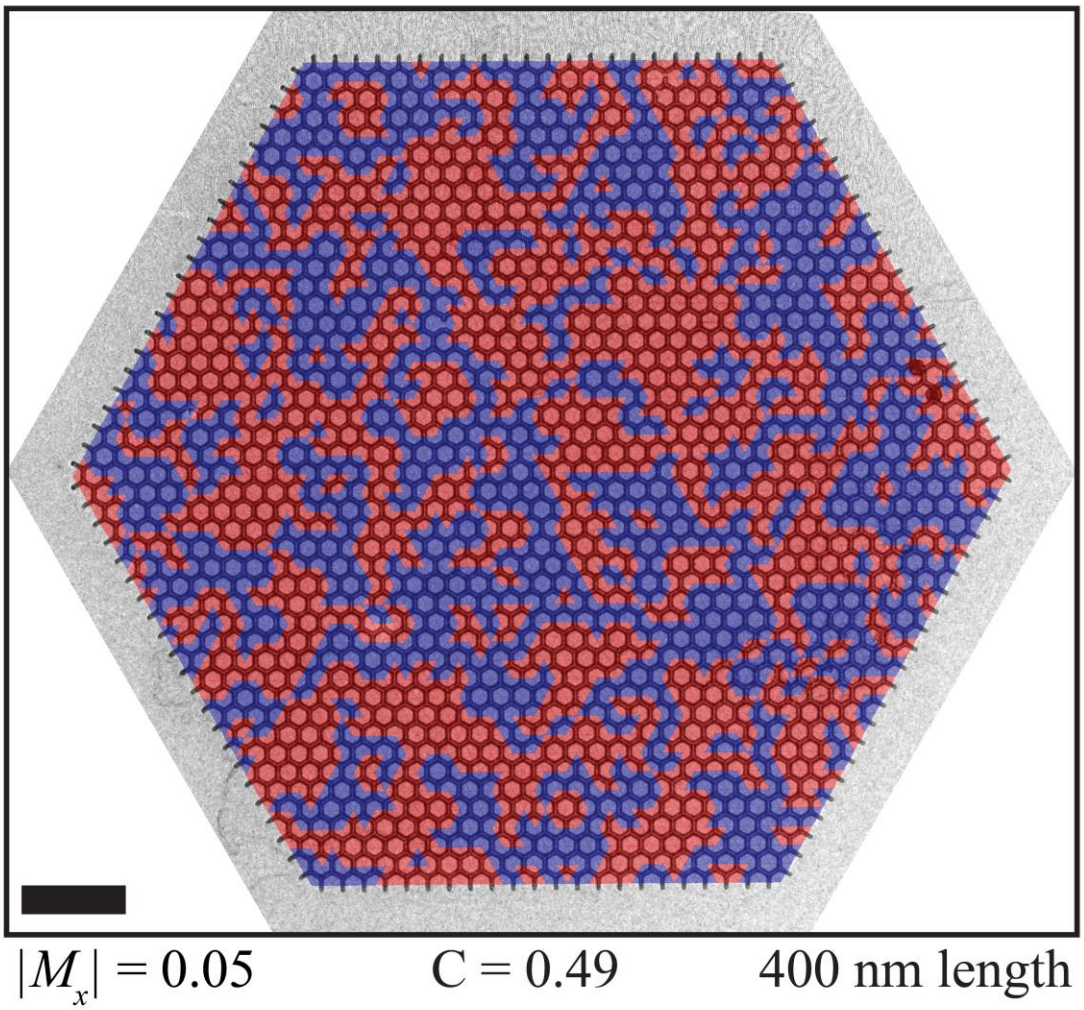

Figure 3.25 **Charge domain map of a 400 nm element length kagome crystal**. Lorentz TEM image with charge domains overlaid showing a characteristic domain pattern for 400 nm element length crystals. This crystal has C = 0.49 and $|M_x|$ = 0.05. The average C for all 400 nm element length crystals studied is given in Figure 3.31. Scale bar is 3 μm.

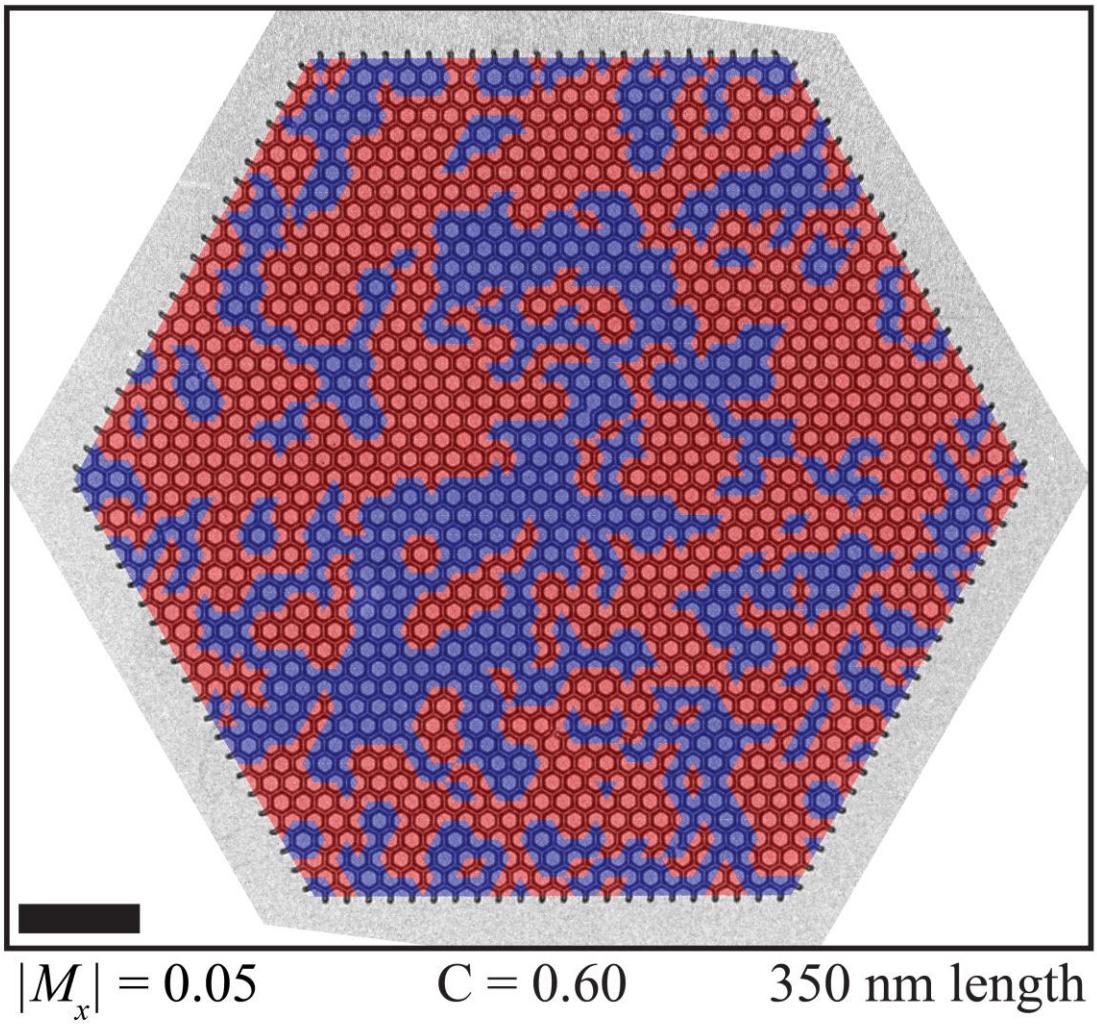

Figure 3.26 **Charge domain map of a 350 nm element length kagome crystal**. Lorentz TEM image with charge domains overlaid showing a characteristic domain pattern for 350 nm element length crystals. This crystal has C = 0.60 and $|M_x|$ = 0.05. The average C for all 350 nm element length crystals studied is given in Figure 3.31. Scale bar is 3 μm.

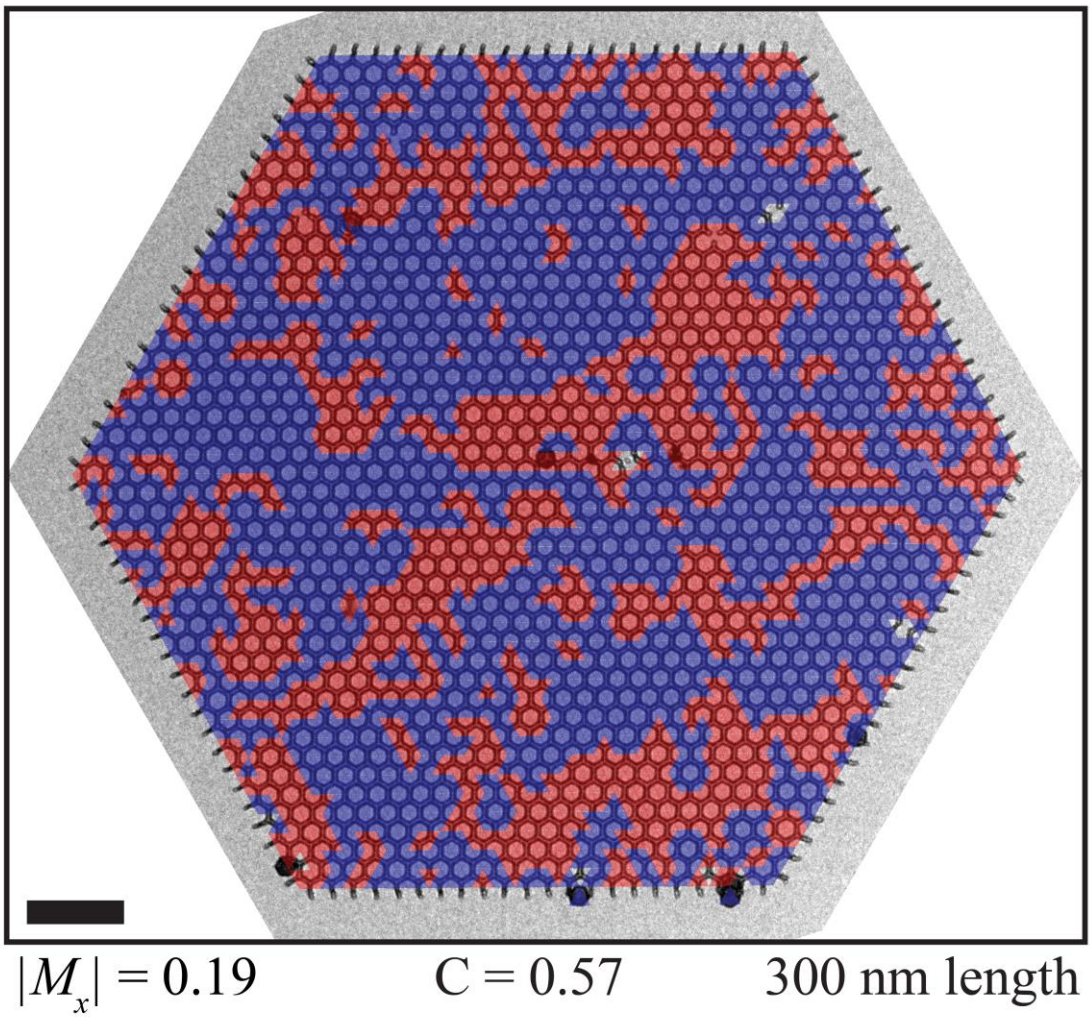

Figure 3.27 **Charge domain map of a 300 nm element length kagome crystal**. Lorentz TEM image with charge domains overlaid showing a characteristic domain pattern for 300 nm element length crystals. This crystal has C = 0.57 and $|M_x|$ = 0.19. The average C for all 300 nm element length crystals studied is given in Figure 3.31. Scale bar is 2 μm.

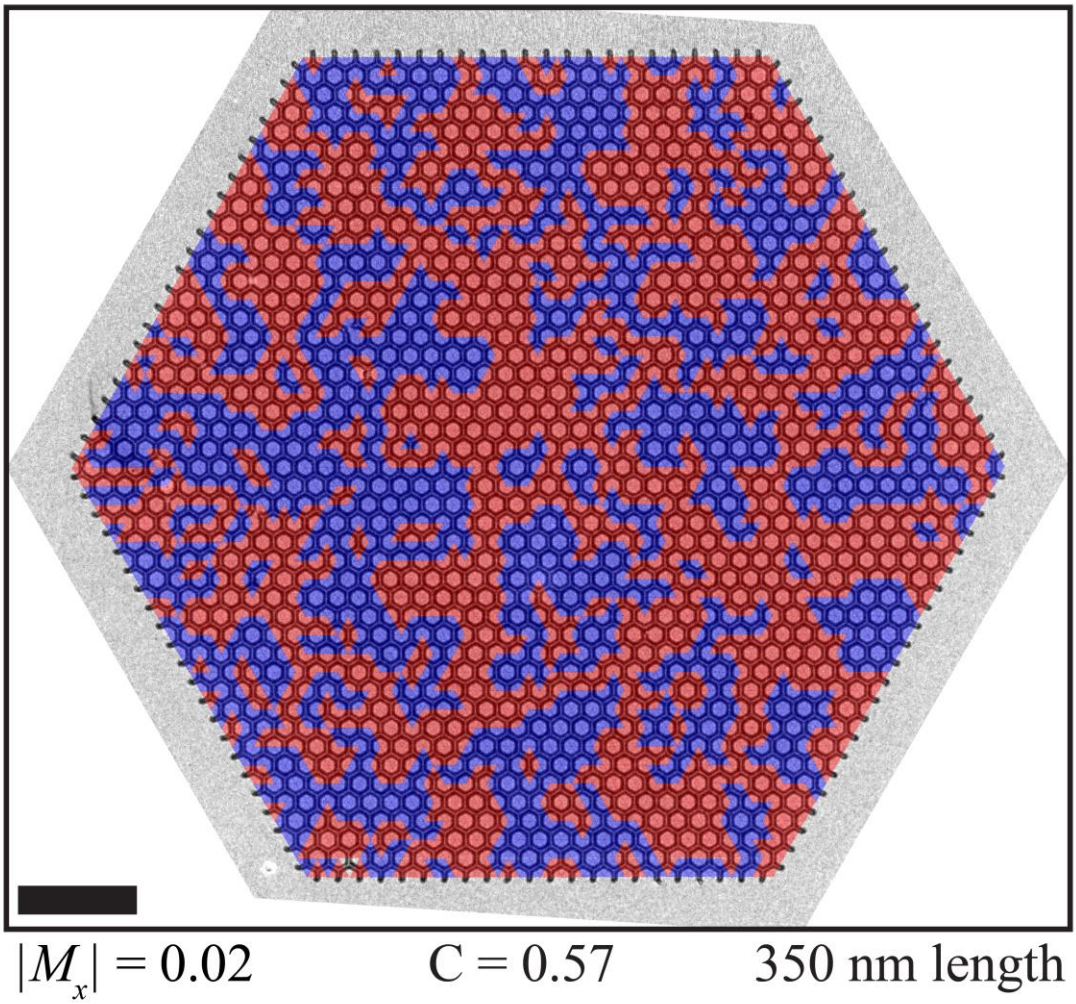

Figure 3.28 **Charge domain map of a 350 nm element length kagome crystal**. Lorentz TEM image with charge domains overlaid showing another characteristic, but well-ordered and demagnetized, domain pattern for 350 nm element length crystals. This crystal has C = 0.57 and $|M_x|$ = 0.02. The average C for all 350 nm element length crystals studied is given in Figure 3.31. Scale bar is 3 μm.

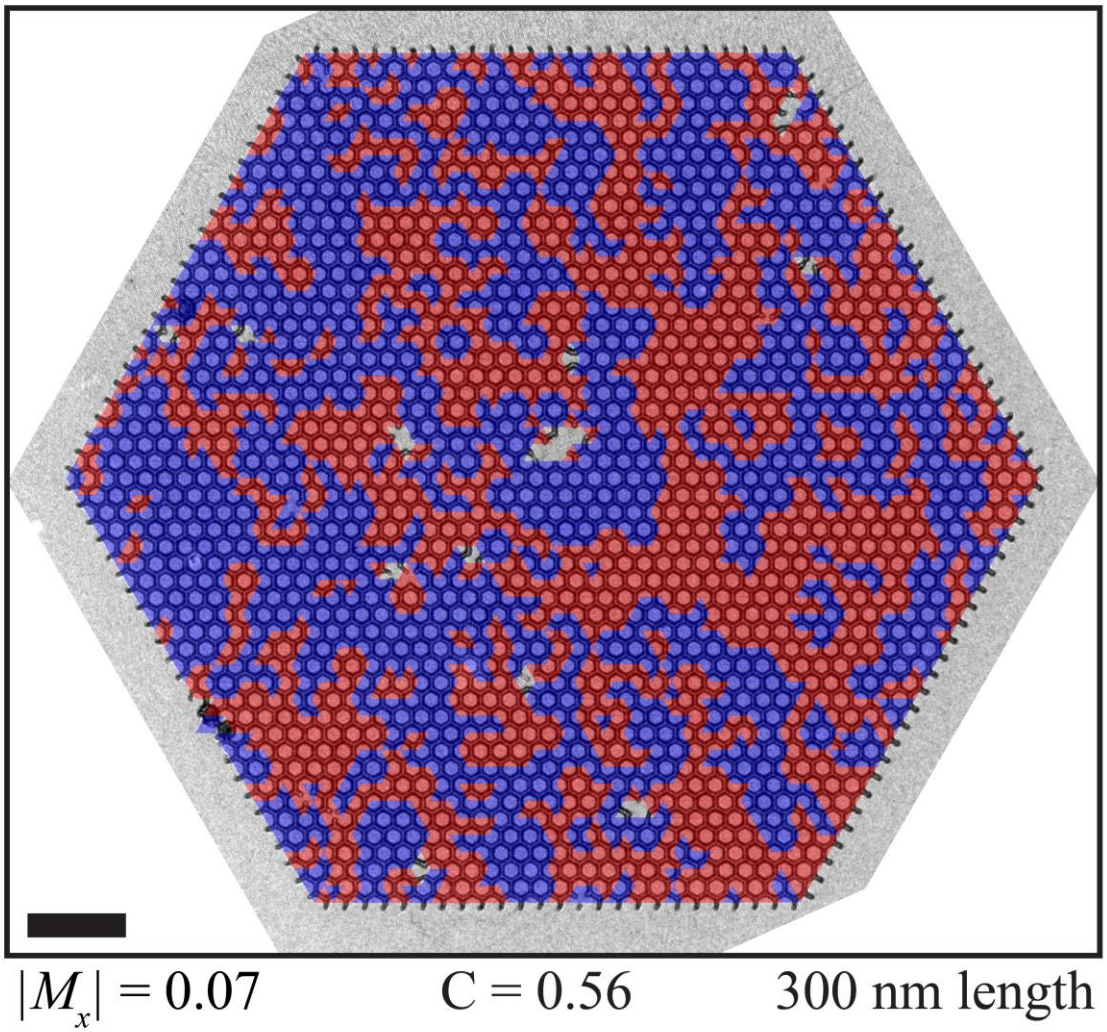

Figure 3.29 **Charge domain map of a 300 nm element length kagome crystal**. Lorentz TEM image with charge domains overlaid showing a characteristic, but well-ordered and demagnetized, domain pattern for 300 nm element length crystals. This crystal has C = 0.56 and $|M_x|$ = 0.07. There are number of vacancy defects in this crystal. The average C for all 300 nm element length crystals studied is given in Figure 3.31. Scale bar is 2 μm.

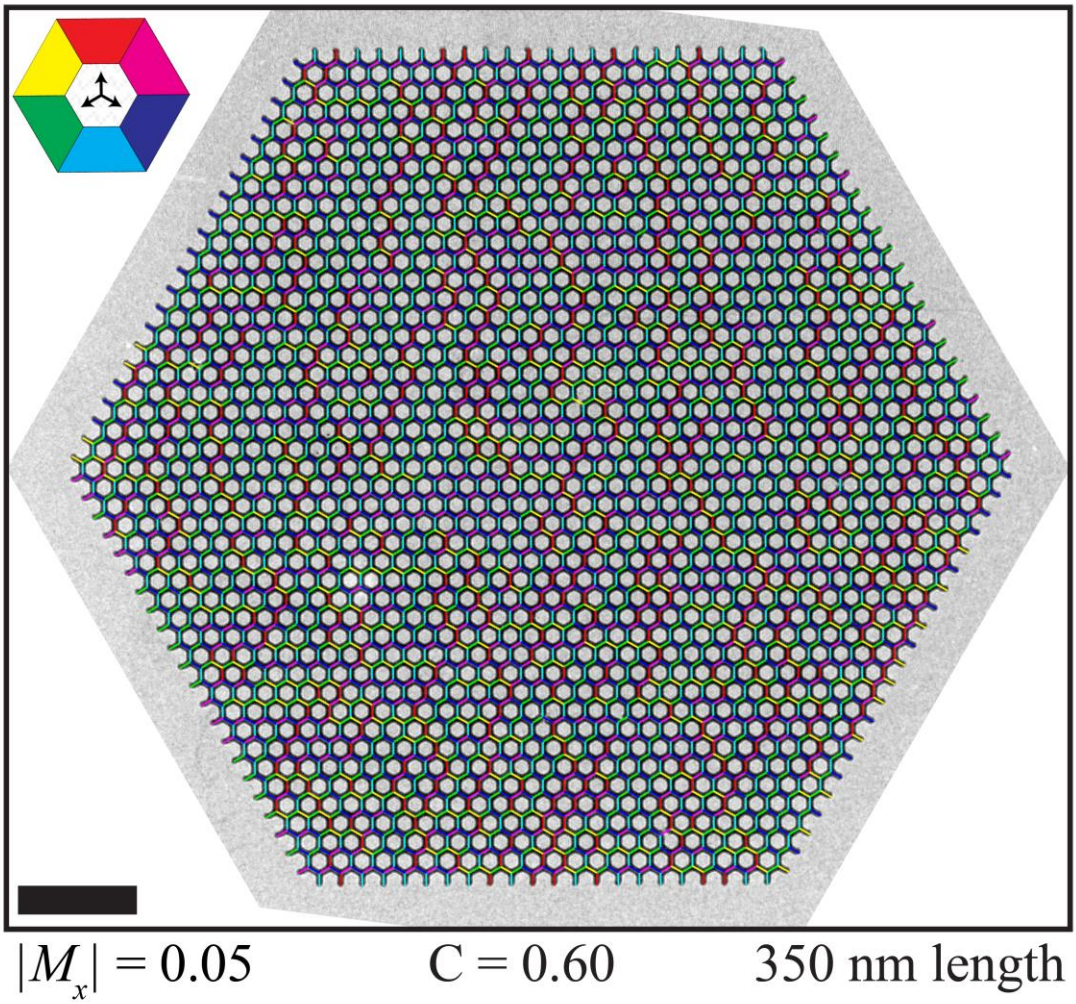

Figure 3.30 **Spin map of 350 nm element length crystal**. Lorentz TEM image with colored lines added indicating the direction of each individual spin for the same 350 nm element length crystal as in Figure 3.26. A color key for the spin directions is given in the upper left corner. The image shows the spins are well demagnetized and randomly arranged, indicative of the kagome ice II state. Scale bar is 3 μm.

We have collected a number of data sets with different length elements, different crystals on one substrate, and different samples made from different depositions of $FePd_3$. When we average our correlation data at each length, a clear trend emerges: shorter lattice constants display greater degrees of charge-ordering. Reducing the element length increases the strength of the magnetostatic interactions between elements and the effective

magnetic charges at the vertices, favoring lower energy configurations. The results of the full data set are shown in Figure 3.31. The shortest element length crystals, 300 nm, do not follow exactly the same trend as the other data points and we discuss this deviation in detail below. For Figure 3.31 we average over 14 crystals for the 500 nm element length, 12 crystals for the 450 nm element length, 25 crystals for the 400 nm element length, 39 crystals for the 350 nm element length, and 6 crystals for the 300 nm element length. The error bars are found by taking the standard deviation of the correlation values at each length and dividing by the square root of the number of crystals studied at that length. Error bars are enlarged in Figure 3.31 for clarity.

It is interesting to compare the 500 nm data point in Figure 3.31 to the *C* values in Figure 3.22 in the previous section. These numbers are approximately equal and within the error bar in Figure 3.31. The 500 nm length medium size geometry crystals tend to show the same behavior as their larger counterparts, which should be expected. This degree of ordering is characteristic of our samples for this element length and it emerges as the kagome crystals increase in size up to a point where they can be confidently averaged over.

The 300 nm element length data point in Figure 3.31 deviates from the trend in the other 4 data points. The point is statistically not in strong disagreement with a "straight" line fit, but we believe this deviation may be due to higher disorder in these crystals and fewer overall successful crystals fabricated (6). This geometry was especially difficult to fabricate because of issues with the liftoff step. The 300 nm length elements have extremely small areas, the "cookies," inside of the connected hexagons that need to be lifted off. The cookies are so small that it is difficult to get a suitable undercut in the bilayer resist, making it difficult for the solvent to efficiently reach underneath the cookie and

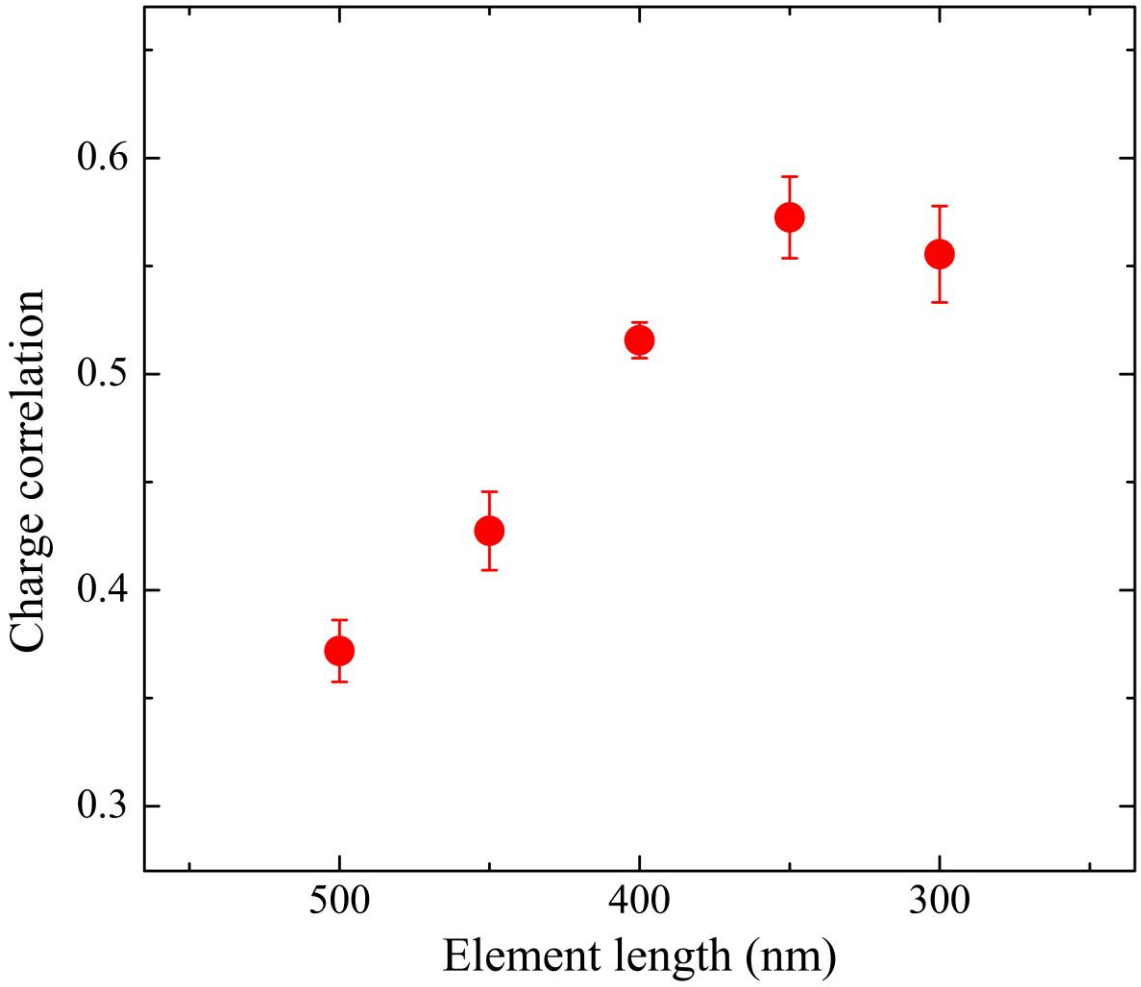

Figure 3.31 **Nearest neighbor charge correlation as a function of element length in the kagome geometry**. Average C of all crystals studied at each element length. Error bars are three times the uncertainty in the average. $|M_x| \leq 0.2$ for all crystals studied.

remove it. Even when there was successful liftoff, the crystals tended to have more disorder than the other sizes.

Figure 3.29 shows a 300 nm element length crystal where more vacancies are present than in the other crystal sizes. There also tended to be a greater size distribution of the widths of individual elements in the 300 nm length crystals. I will discuss the width variation further in the next section on Monte Carlo modeling, but recall that Budrikis *et al*. showed that disorder in ASI systems in known to have similar overall effects even when

it comes from different sources [121]. The disorder in this case is affecting how well the system can find a charge-ordered state. Both increased numbers of vacancies and a large spread in widths of individual elements contribute to the lower than anticipated charge-ordering for this size element crystals. Based on the data from the 500 nm to the 350 nm element length crystals, these crystals should find highly charge-ordered states. Some of the 300 nm element length crystal do order very well, but others have higher disorder and there are only 6 successful crystals at this length to study, which results in an average charge correlation that deviates from the trend for the other element lengths.

We can also calculate higher order charge correlations between vertices. In Figure 3.31, we quantified the charge-ordering using only nearest neighbors, but we can also calculate similar values for further charge neighbors. When doing this, most of the higher order correlations can be calculated in exactly the same way as the parameter *C* we have already defined, but it is important to realize that for some specific neighbor distances, this would yield a parameter of opposite sign from our convention. We correct for this by defining the correlation to be 1 if the pair of charges fits into the same charge domain and -1 if the charges belong to different domains.

The results of these calculations for up to 7 nearest neighbors are shown in Figure 3.32. It is also interesting to remember that a correlation value of 0 represents a random arrangement of spins and charges, so even for the furthest charge neighbors we studied in our samples, the 7^{th} nearest neighbors, the correlation is generally greater than a random configuration. This is definitive evidence of long-range order, long-range correlation, and long-range interactions in this system. The magnetic charges at the vertices interact through

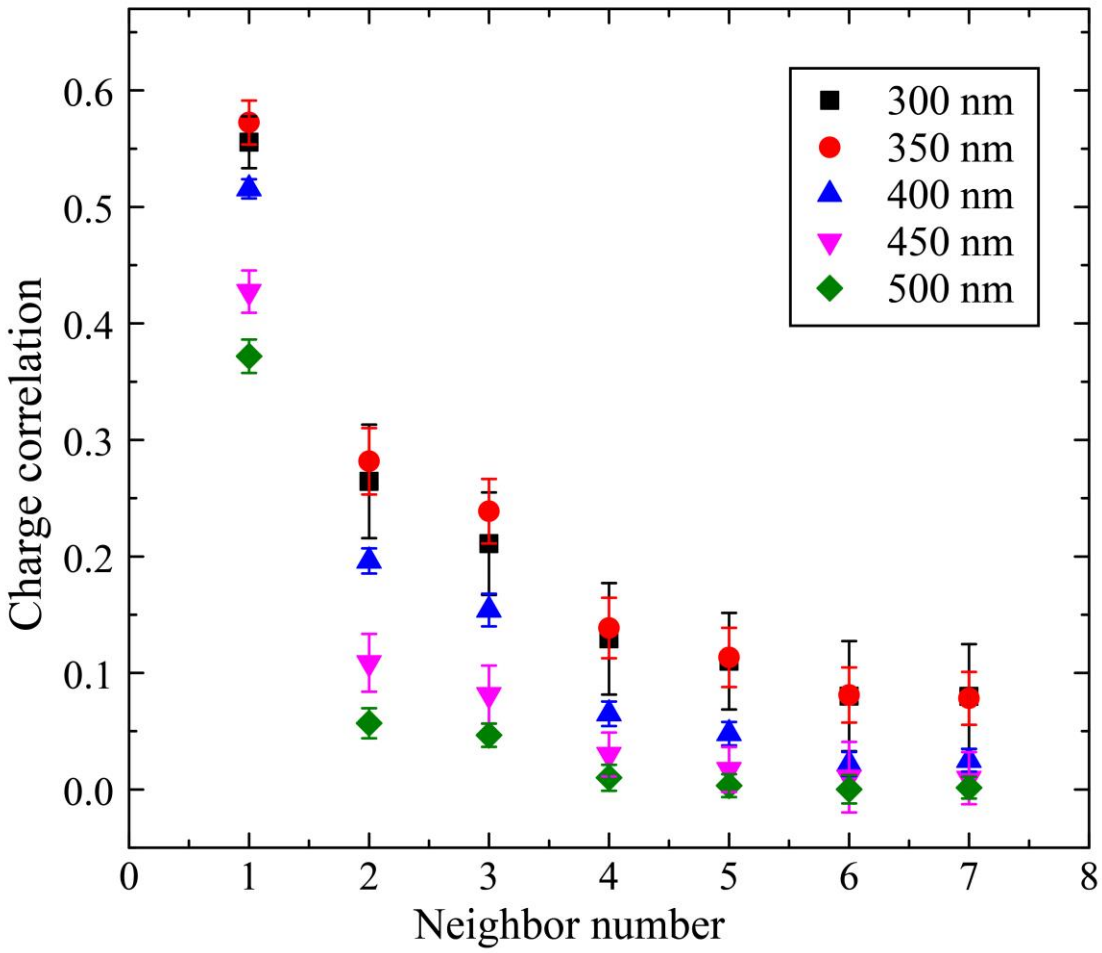

Figure 3.32 **Higher order charge correlations**. Average charge correlation over all crystals studied as a function of neighbor number for each element length. Error bars are three times the uncertainty in the average for clarity.

a Coulomb term with similar long-range effects and fall-off as dipolar interactions, giving this system long-range order highlighted by the data in Figure 3.32.

Overall, the charge-ordering data presented here show the highest level of ordering yet achieved to date in an experimental kagome ASI system. Our results still deviate from the predicted perfectly ordered behavior and we believe this is due mainly to the disorder in our system. In the next section, I describe by kinetic Monte Carlo simulations that I used to help understand what role disorder may play in charge-ordering. With the prospect that

this disorder can be overcome, we believe these results pave the way for future investigations of perfectly charge-ordered and even spin-ordered states in the kagome geometry.

3.4 Monte Carlo modeling of large kagome crystals

In this section we report our results on kinetic Monte Carlo simulations modeling the data presented in Section 3.3 on charge-ordering of large kagome crystals as a function of element length. While we have experimentally achieved the highest yet degree of charge-ordering so far reported in the kagome geometry, our results are still far from the predicted perfectly charge-ordered state or spin ordered states. In order to better understand our system and its deviation from theoretical predictions, we use Monte Carlo methods to solidify the model for our system. The framework for the kinetic Monte Carlo technique is given in Section 2.5.

The major difference between my kagome Monte Carlo simulations and the square geometry Monte Carlo simulations presented in Section 2.5.2 is how the ΔE term is calculated. We found that it is more effective to calculate the change in energy from magnetic Coulomb interactions, similar to other reports on the kagome geometry [59,80]. This interaction is given by

$$V_{ij} = \begin{cases} \frac{\mu_0}{4\pi}\frac{Q_i Q_j}{r_{ij}}, i \neq j \\ \frac{\mu_0}{4\pi} Q_i^2, i = j \end{cases} \quad (3.3)$$

where Q_i is the total charge at each vertex, given by the sum of individual dumbbells of charge and r_{ij} is the distance between the two vertices. Each half-dumbbell of charge at a vertex is given by whether it is a $+Q$ or $-Q$ charge times the magnetic moment of the

element divided by its length [191]. Since the magnetic moment is the magnetization times the volume, the total charge at each vertex simplifies to

$$Q_i = \sum_{j=1}^{3} q_j \frac{\mu}{l} = \sum_{j=1}^{3} q_j MA \tag{3.4}$$

where M is the magnetization of the material and A is the cross sectional area of an element. When Q_i and Q_j are of opposite charge, this lowers the total energy which is favorable. The bottom term in Equation 3.3 is the self-energy of each magnetic charge. Since it is proportional to the charge squared, it is therefore highly energetically unfavorable to have a $\pm 3Q$ charge, which explains why we don't see any of these in our experimental systems.

In our experimental studies, we see the degree of charge-ordering improve as the element length decreases. In our Monte Carlo simulations, this factors into the ΔE term as it changes the magnitude of r_{ij}. We also need to take into account the widths of the individual magnets in our simulations. In our experimental samples we vary the length of the elements and as we do, the widths of the magnets tend to vary as well due to proximity effects during lithography. When the vertices connecting three elements are fabricated closer and closer together, this directly affects how the resist is exposed in the regions between the vertices. When other elements are fabricated closer to a given element, it takes a smaller dose to expose the same region. We attempted to adjust the dose to account for this, but we still find that shorter element lengths tend to result in wider magnetic elements. As such, we vary the element width with length in our simulations as $w = -0.2l + 190$. This is shown as the line in Figure 3.33 and is representative of the widths of our real elements. The experimental data points in Figure 3.33 are measured from high resolution TEM images like the one in Figure 2.18(d). We also use these high resolution images to measure

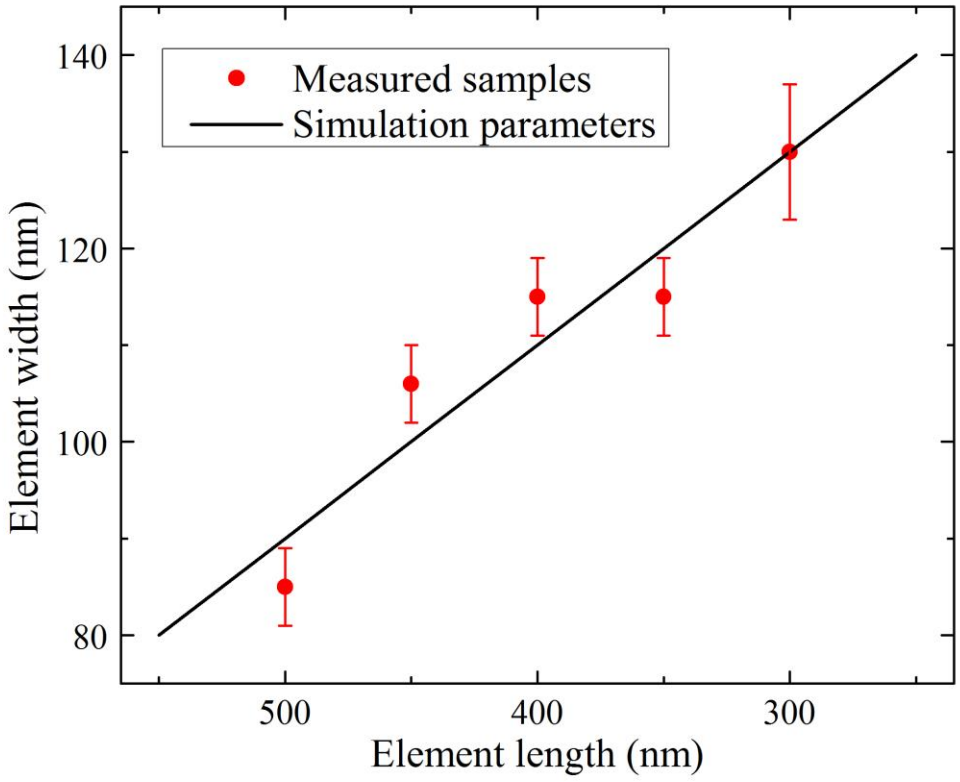

Figure 3.33 **Simulation element sizes**. Element width vs. element length measured from TEM images of our samples and a line showing correlated length and width input parameters in our simulations. Error bars are one standard deviation of measured widths. Note the 300 nm length elements have a larger spread in widths than the other lengths.

the characteristic spread in widths of our individual elements, given by the error bars in Figure 3.33.

With no disorder, the computational simulations produce perfect $C = 1$ ordering every time, but we find that when we include a small amount of random disorder we are able to reproduce our experimental results quite well. There have been a number of studies focused on how disorder affects ground state formation in artificially frustrated systems [118–121]. In our samples, there is a disordered spread in the width of the magnets due to lithography artifacts, arising from slight local variations in electron dose due to shot noise.

A 4 nm standard deviation of widths is representative of the real spread in element widths in our 500-350 nm length samples from our microscopy observations (see Figure 3.33). The 300 nm length samples are very difficult to fabricate and tend to have more edge roughness and a larger spread in element widths, closer to 7 nm. There were also fewer successful crystals made at this size. This explains the deviation of the 300 nm data point from the model.

We run the simulations with the same number of elements (~5700) and same boundary conditions as our real samples. We use a magnetization of 200 kA m^{-1} and evolve the simulation for an average of 200 flips per spin. We also repeat each point 75 times with different randomly distributed disorder and average over the results. As Figure 3.34 shows, when we assign each magnet a width randomly sampled from a Gaussian distribution with a standard deviation (σ_w) of 4 nm, we find very good agreement with our experimental data. By increasing the standard deviation of the distribution to 5 nm, we are able to reproduce the 300 nm data point and also demonstrate the same upward trend in the correlation as a function of decreasing element length. The results of these simulations are presented in Figure 3.34.

The width disorder in our samples enters the simulation in two ways. As was discussed in Section 2.5, the intrinsic energy barrier E_0 depends on both the volume of the element and D_Z which are both strong functions of the width of the magnet. Disorder also enters through the ΔE term, as the magnitude of the half-dumbbell of charge depends on the cross sectional area of the individual element which is directly proportional to its width. Both of these effects cause slight variations in the energy landscape of the system which directly affect the relative flip rates for different individual spins.

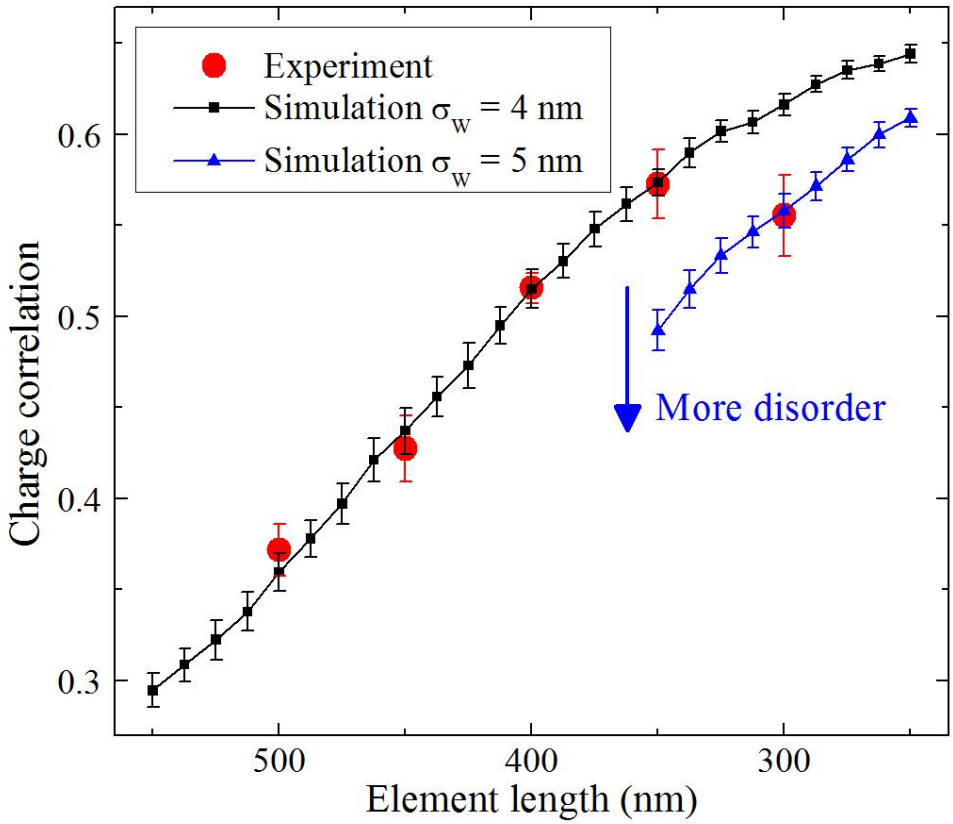

Figure 3.34 **Monte Carlo results**. Charge correlation as a function of element length in the kagome geometry for our experimental results and kinetic Monte Carlo simulations including two levels of width disorder. Experimental error bars are three times the uncertainty in the average, simulation error bars are two times the uncertainty in the average for clarity. The Monte Carlo results reproduce the experimental data extremely well.

Figure 3.34 shows excellent agreement between theory and experiment for quantitative charge-ordering in our system as a function of element length and corresponding element width. We can also create charge-domain maps from our simulations, just like the experimental ones found in Figures 3.23-3.29. Figures 3.35-3.39 show characteristic charge domains for our kinetic Monte Carlo simulations as a function of element length and width. These are final frames from one of the 75 simulations ran at each data point. In these images, it is clear that the average size of the domains increases

as the element length decreases, just as we observe in our experimental data. These domain outputs look remarkably similar to the experimental ones, strongly suggesting that we have a good numerical model of our system.

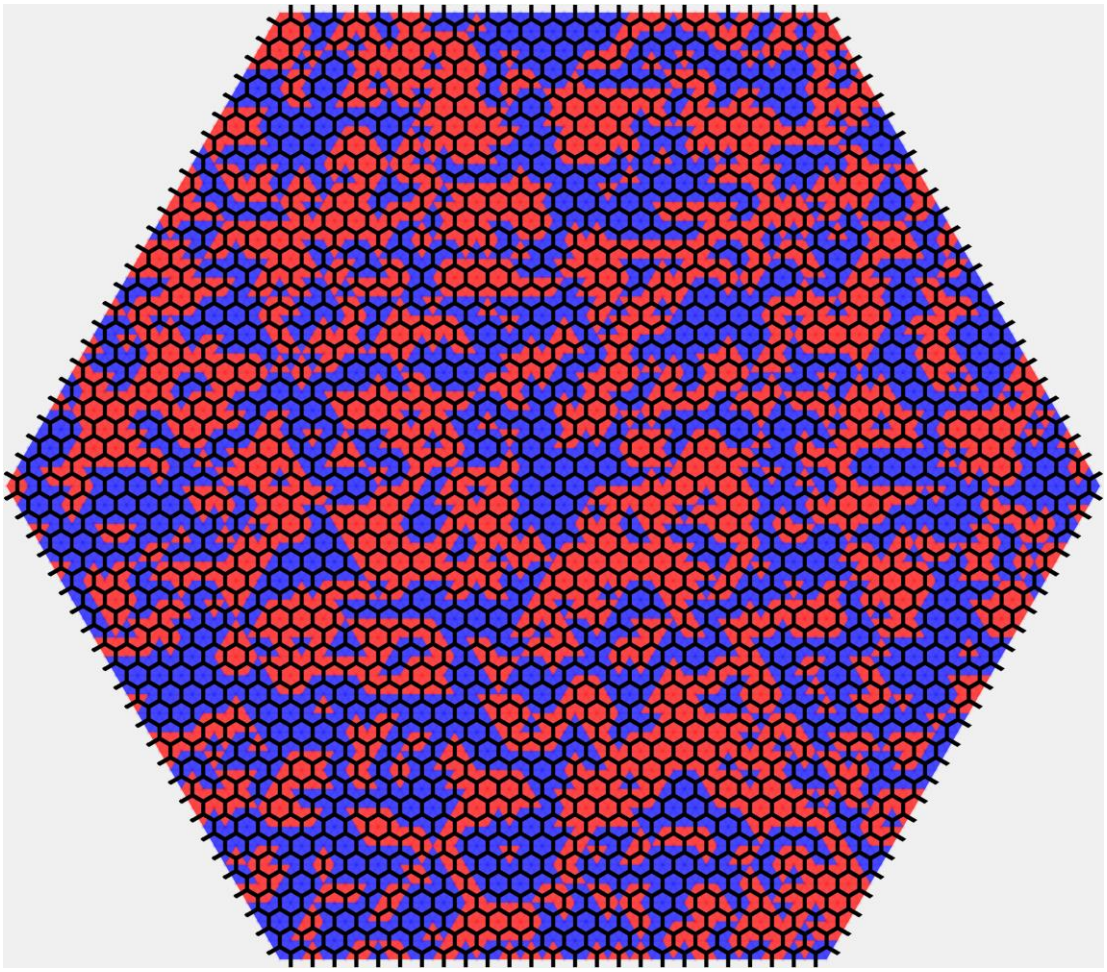

Figure 3.35 **500 nm element length simulation charge domains output**. Characteristic charge domain pattern for 500 nm element length kinetic Monte Carlo simulation of the kagome geometry. This run has an element width of 90 nm and a spread in widths σ_w = 4 nm, which match measured experimental samples. The charge correlation C = 0.37, which is approximately the average over both all experimental crystals and all simulation results.

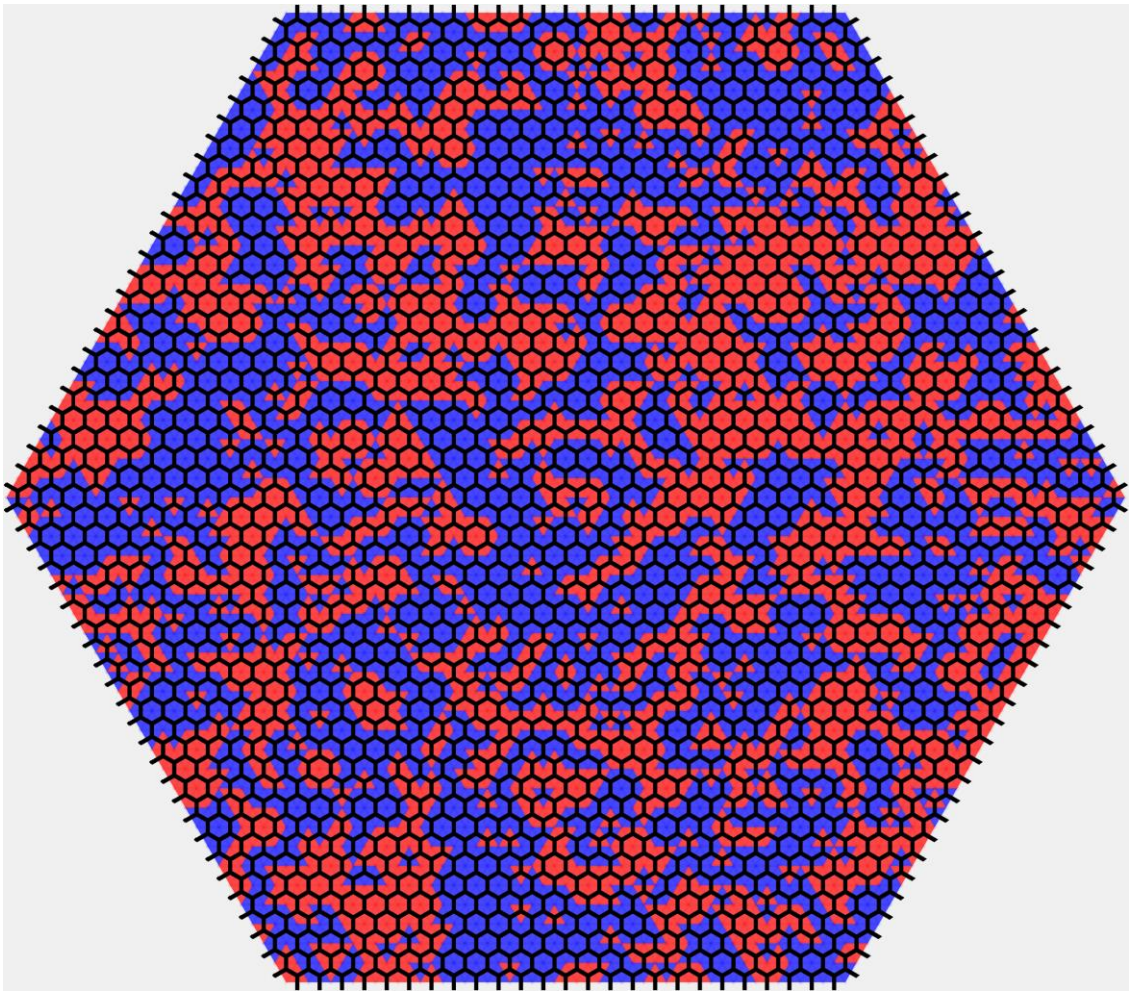

Figure 3.36 **450 nm element length simulation charge domains output**. Characteristic charge domain pattern for 450 nm element length kinetic Monte Carlo simulation of the kagome geometry. This run has an element width of 100 nm and a spread in widths σ_w = 4 nm, which match measured experimental samples. The charge correlation C = 0.42, which is approximately the average over both all experimental crystals and all simulation results.

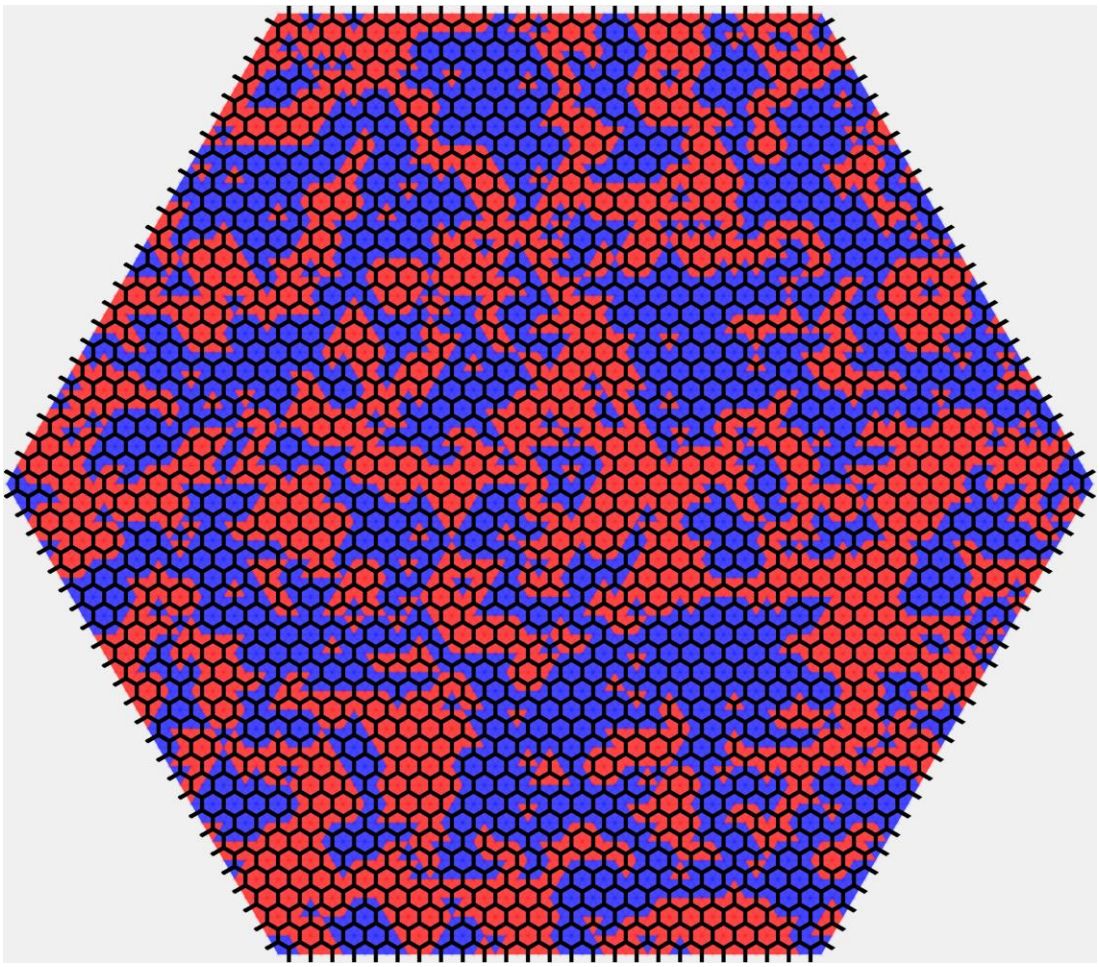

Figure 3.37 **400 nm element length simulation charge domains output**. Characteristic charge domain pattern for 400 nm element length kinetic Monte Carlo simulation of the kagome geometry. This run has an element width of 110 nm and a spread in widths σ_w = 4 nm, which match measured experimental samples. The charge correlation C = 0.50, which is approximately the average over both all experimental crystals and all simulation results.

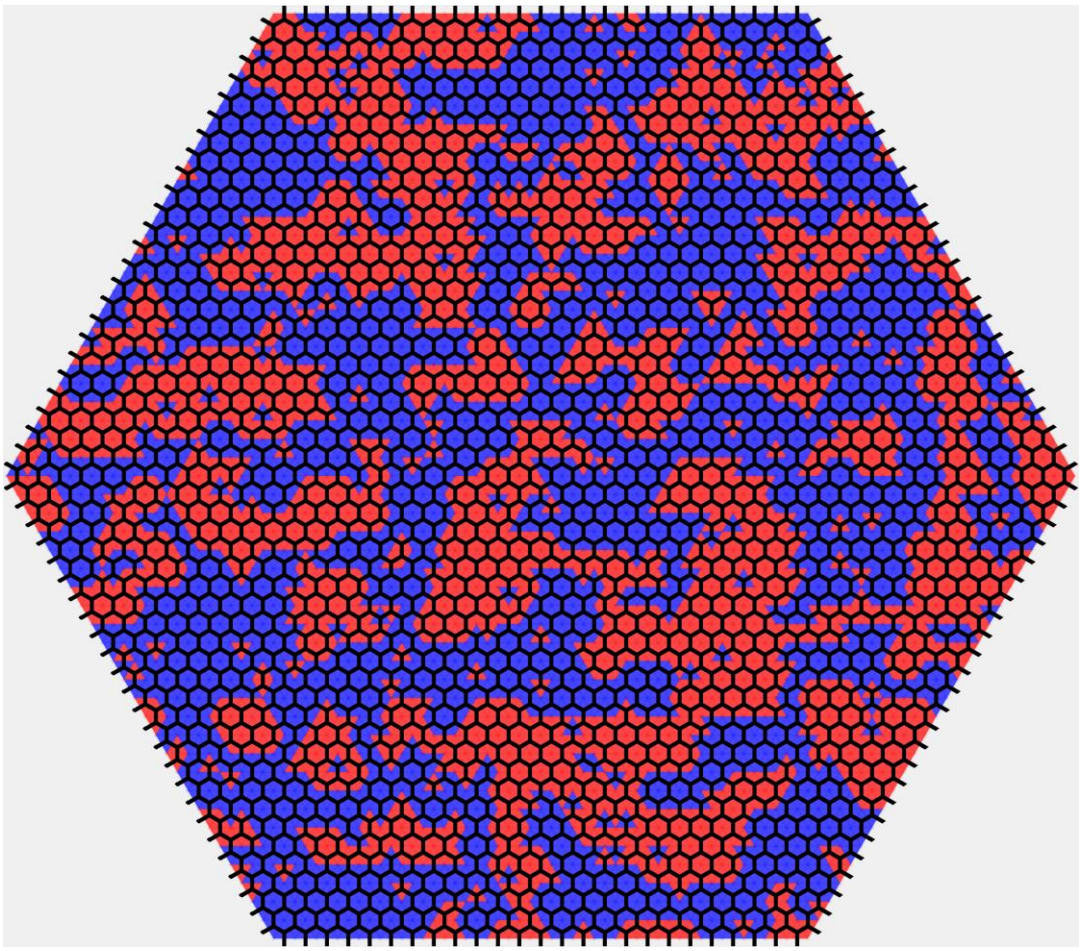

Figure 3.38 **350 nm element length simulation charge domains output**. Characteristic charge domain pattern for 350 nm element length kinetic Monte Carlo simulation of the kagome geometry. This run has an element width of 120 nm and a spread in widths σ_w = 4 nm, which match measured experimental samples. The charge correlation C = 0.58, which is approximately the average over both all experimental crystals and all simulation results.

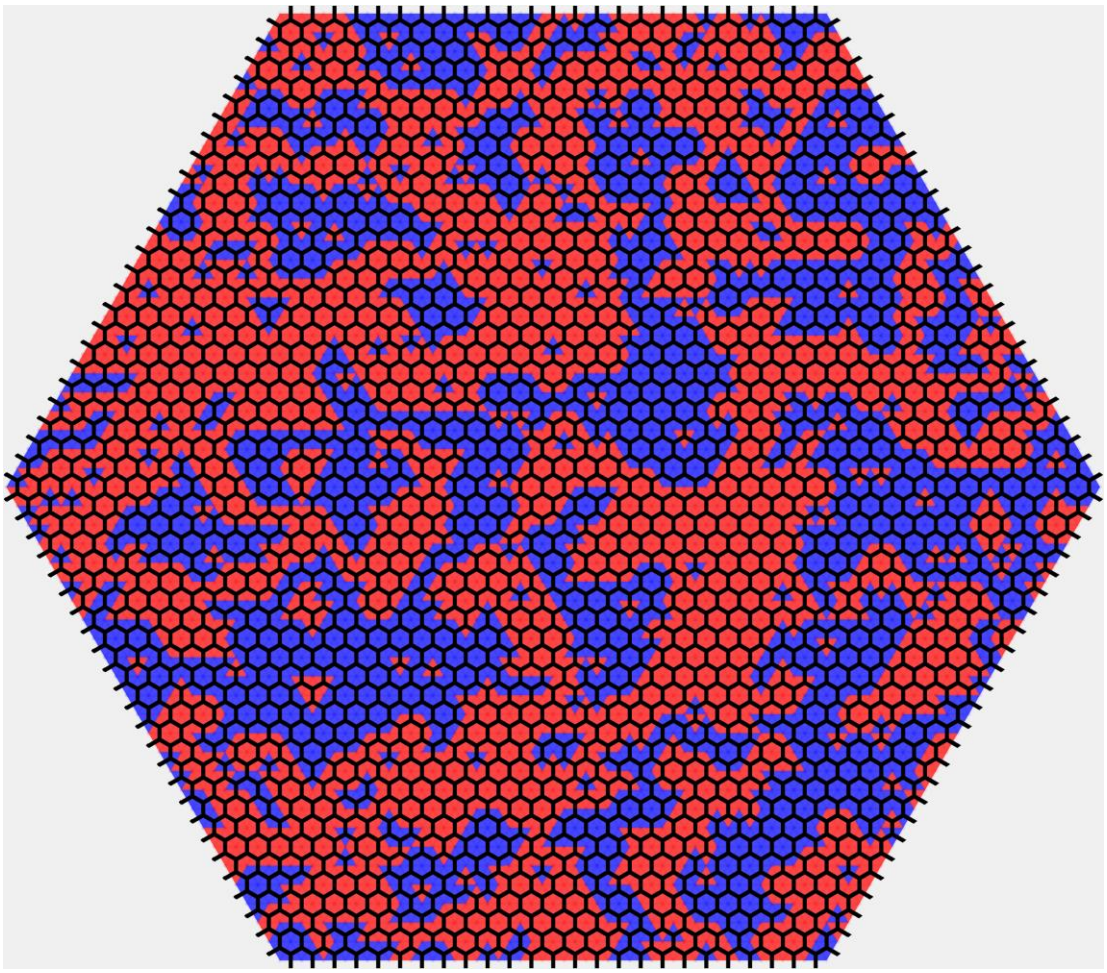

Figure 3.39 **300 nm element length simulation charge domains output**. Characteristic charge domain pattern for 300 nm element length kinetic Monte Carlo simulation of the kagome geometry. This run has an element width of 130 nm and a spread in widths σ_w = 5 nm, which match measured experimental samples. The charge correlation C = 0.55, which is approximately the average over both all experimental crystals and all simulation results.

Having found a good model for our kagome geometry samples that can reproduce the experimental data very well, it is interesting to compare the results of the kagome simulations with our square geometry simulations discussed in Section 2.5.2. In the square geometry Monte Carlo work, we find that the same choice of parameters and approximately same amount of width disorder (4 nm) as the kagome simulations still produces very good long-range ordering with large domains of ground-state type I vertices in the square lattice,

similar to the experiments shown in Figures 2.35, 2.37, and 2.38. This shows that the critical slowing down of spin dynamics in the kagome geometry [127] may be more susceptible to disorder, further preventing perfect long-range ordering in this system. We include disorder as a distribution of widths in our system which is reasonable given our high-resolution TEM observations. The actual source of disorder may be more than simply width disorder; as I discussed how different types of disorder have similar effects in ASI systems [121]. Our numerical results still capture the fundamental ideas and the overall behavior of the system extremely well and give a plausible explanation for the observed deviations from theoretical predictions [125,127].

3.5 Conclusions and future work

In the previous sections of this chapter, we presented results from studied on the charge-ordering properties of kagome geometry samples from the smallest possible crystals up to large arrays with thousands of individual elements. We model our large kagome systems using a kinetic Monte Carlo technique, and we are able to reproduce our experimental data extremely well by including a small amount of disorder in the simulations. In the smallest crystals, we find that edge effects play an important role, as head-to-tail flux-closure arrangements of spins compete with higher order interactions that favor charge-ordering. This observation is consistent with the emergent frustration of the system. The smallest crystals have a limited number of possible states and their finite size causes edge effects to dominate the effects of the frustration in some cases. For larger crystals, enumerating all the possible states and ground states quickly becomes intractable,

the edge effects play a negligible role, and the frustration is much more apparent. We were able to observe perfectly charge-ordered crystals up to 7 hexagons in size and demonstrate extremely reproducible behavior in one, two, and three-hexagon crystals.

In Section 3.3, we presented charge-ordering behavior as a function of element length for large kagome crystals. Here, we realize the highest yet degree of charge-ordering reported in the kagome ASI geometry by taking advantage of the benefits of $FePd_3$ as a material for studying thermally active ASI systems. We find that as the length of individual elements decreases, the degree of charge-ordering and the size of the charge domains increase due to stronger magnetostatic interactions between neighboring effective magnetic charges. In Section 3.4, we model the results presented in Section 3.3 and attribute the samples' deviations from predicted perfectly ordered states to disorder in the system due to lithography artifacts and fabrication. Future work may aim to achieve even further ordering than presented here. Small kagome crystals were only fabricated with 500 nm element lengths and further investigations could yield new insights and larger perfectly ordered crystals by using shorter lattice spacings.

A two stage ordering process is predicted for the kagome geometry, characterized by two plateaus in the entropy of the system and two corresponding phase transitions [98]. The first plateau is the magnetic charge-ordered state, followed by a further reduction of entropy into a spin-ordered state. Presumably, this state could be achieved by reducing the disorder in our samples. Careful studies of resist thickness, dose, and accelerating voltage in the lithography step will help reduce the spread in widths of individual elements. A move to higher energy lithography processes, the use of a smaller aperture to reduce beam current and slow beam movement, or a study of different types of resist could all improve the

spread in widths we observe in our samples. Also, Ar ion milling is an inherently disordered technique and new ways to transfer the ASI pattern into the $FePd_3$ films could be explored. In light of the success in easily achieving the square geometry ground state and our very well charge-ordered kagome lattices, we believe $FePd_3$ creates a promising pathway for reaching the first plateau with perfect charge-ordering, paving the way for future studies probing fundamental kinetics of the second phase transition. The experimental observation of this phase transition is of great interest and could shed light on analogous phase transitions in real frustrated materials like the pyrochlore spin ices.

Another major route to improve our results would be through improving the quality of $FePd_3$ films we fabricate. The process I developed using a modular target was excellent as a proof of concept and yielded very nice and promising results. There is still a wide range of film qualities produced and throughout all my work it is evident that the starting quality of the films directly impacted both the quality of the Lorentz contrast images and the quality of the final data. Simply put, better thin films result in better long-range ordering. Future work could involve exploring ways to reduce the surface roughness and grain sizes in the films and how to fabricate only the highest quality films more consistently. Using an alloy $FePd_3$ alloy target could likely create more uniform films and depositing the films in the annealing chamber could likely help as well. I had to break vacuum between deposition and annealing, and I did not do controlled studies of UHV annealing protocols for the $FePd_3$ films. Studying the effects of different annealing temperatures, annealing times, quenching, and annealing during deposition could all impact the final favorable magnetic properties of the films. Different underlayers and underlayer thicknesses could be explored as well.

Our square geometry samples find highly ordered states with large regions of ground-state order. Under the same annealing conditions, the kagome geometry samples do not show the same level of order, likely due to the kagome geometry being more frustrated. The local vertices of the kagome geometry are more degenerate than the local vertices in the square geometry and the kagome geometry needs higher order interactions to realize the charge-ordered and spin-ordered states. With the exciting results presented here, and the prospect that the disorder in our samples can be overcome and the quality of our films can be improved, our system shows great promise as a model system for future studies. In particular, it may allow for even higher degrees of ordering and new explorations into more sensitive order parameters, closer to the true kagome ground state.

Chapter 4 Topological frustration of artificial spin ice

4.1 Introduction

In this chapter, we introduce controlled topological defects into square ASI lattices and directly observe the resulting spin configurations. This work is quite distinct from the experiments and theory discussed in previous chapters and the existing body of work on ASI systems in general. As such, this topic deserves its own thorough introduction and motivation. I will first discuss the idea of a dislocation as a topological defect and the diverse and interesting phenomena associated with these defects in different systems. I will then survey previous studies related to topological frustration in a variety of systems and define exactly how we interpret and apply this idea to ASI. With some motivation from crystalline systems such as the pyrochlore spin ices, I will describe how we introduce topological defects into our ASI samples, the fabrication details of our system, and all of the different defect geometries we study. I will then present extensive results from our system and show the diverse behaviors we observe specifically due to topological frustration. Finally, I develop a theoretical framework for topological frustration, simulations using the kinetic Monte Carlo technique, and I use the results to make predictions about the role of topological frustration in a variety of ordered systems.

An ASI approach to topological frustration gives detailed experimental information at the single-spin level, which is generally not accessible in real materials. This opens the possibility for direct study of specific crystalline imperfections. Topological defects, such as dislocations [193], have an unavoidable presence in crystal lattices although they are often neglected in experimental and theoretical investigations due to their complexity and

the difficulty in resolving their precise atomic arrangements. The results we present in this chapter demonstrate the non-trivial nature of topological defects, with implications for many real systems in which a typical density of dislocations could fully frustrate an otherwise unfrustrated system. We interpret our results in the context of unanswered questions about the pyrochlore spin ices [13,27,51,55,57], and point to some other materials where the presence of topological defects may also play a strong role. It is striking to note how little the role of dislocations have been considered in the extensive literature on ordered systems.

Dislocations are topological defects [193] ubiquitous in crystalline materials and can evoke a diverse range of phenomena across vastly different systems. Examples of these phenomena include: theoretically predicted ferromagnetic dislocations in an antiferromagnetic lattice [194–197], which have recently been experimentally observed in antiferromagnetic NiO [198]; one-dimensional fermionic excitations in topological insulators [199]; plasticity in metallic alloys [200]; and a recent report on confined structural states at dislocations in an FeMn alloy [201]. The topological nature of dislocations means their presence can be measured far away from the actual defect site by following a path around the dislocation core, resulting in an uncompensated lattice constant and requiring an extra vector to complete the loop, known at the Burgers vector [17]. Figure 4.1 shows a schematic of a topological defect in a 3D crystal and Figure 4.2 presents a defect in a 2D lattice and gives a definition of the Burgers vector. While non-topological defects such as vacancies or substitutions can impart useful and favorable properties for crystals in, for example, semiconductor engineering, the presence of such a defect is generally unknown in parts of the material far removed from it. However, the non-trivial

presence of even a single dislocation can produce long-range topological effects that permeate the crystal. Typical dislocation densities in metals are on the order of 10^5-10^{12} cm^{-2}, while in ceramics they can be much lower, around 10^4-10^6 cm^{-2} [202]. For this reason, dislocations generally receive less attention in ceramic materials.

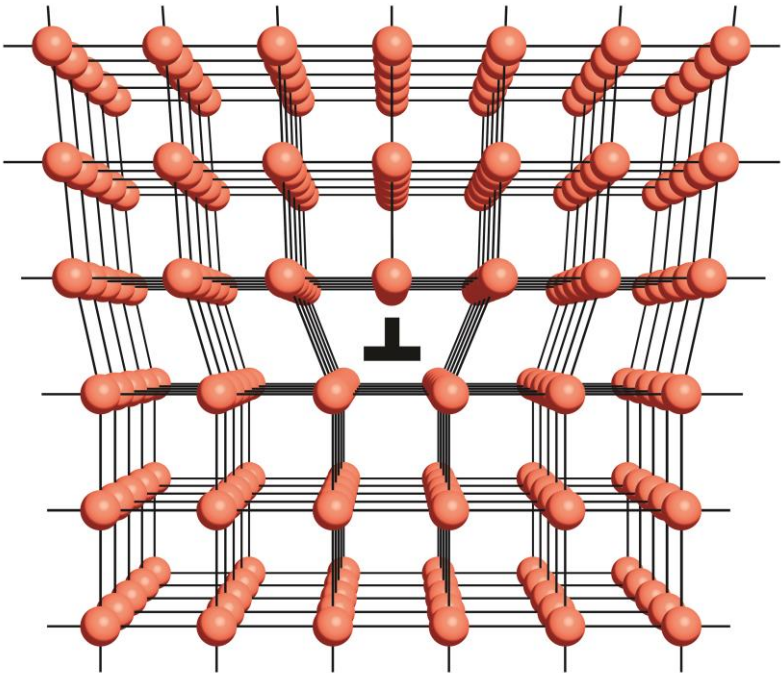

Figure 4.1 **Dislocation in a 3D crystal**. Schematic of an edge dislocation in a 3D crystal. At the dislocation point the number of lattice constants changes from 6 to 7 and an extra plane of atoms appears. The conventional "T" notation for a dislocation is also shown. Figure adapted from Ref. [203].

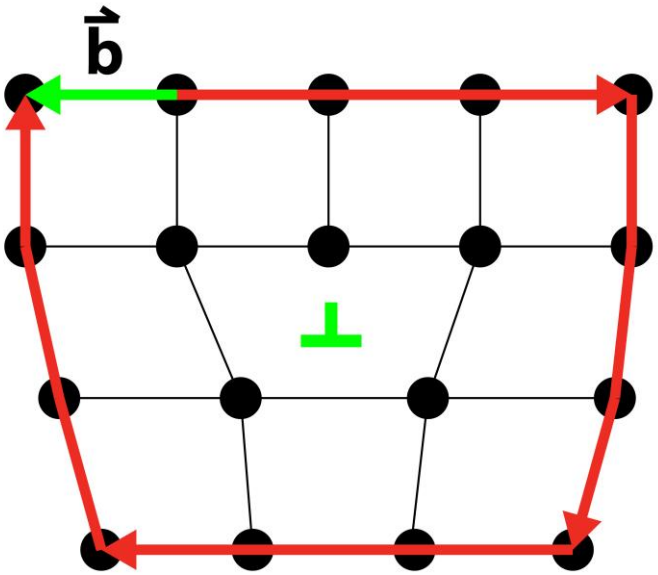

Figure 4.2 **Topological defect in a 2D crystal**. An edge dislocation like the one in Figure 4.1 projected onto a 2D crystal. The Burgers circuit and Burgers vector ($\vec{b}$) are also indicated. The dislocation is a topological defect because any loop around the dislocation, no matter how big, with an equal number of lattice constants in each direction cannot be completed without adding an additional vector, known as the Burgers vector.

An important consideration in the relationship between structure and frustration that has received relatively little attention is the presence of topological structural defects. Topological influences on frustration have been studied in protein folding [204–208], in the carbon bonding of graphene nanoflakes [209,210], in nematic liquid crystals [211], and in spin configurations at the Cr(001) surface [212,213]. However, the level of control and direct characterization in these systems leaves numerous unanswered questions. Here, we present a controlled and direct study of the topological frustration, specifically due to topological defects, in otherwise ideal crystals. We introduce controlled topological defects in the form of edge dislocations into our thermally active square ASI systems and directly observe the resulting frustrated spin configurations upon annealing. The individual plaquettes of the system are locally unfrustrated, but the whole structure remains frustrated due to its topology, leading to a high number of low lying energetic states. Figure 4.3

presents a schematic representation of topological frustration as compared to other well-studied sources of frustration. Geometric frustration has been discussed at length thus far in this dissertation, and geometric considerations are the source of the frustration in both the square and kagome ASI lattices as well as in the pyrochlore spin ices. Disorder frustration is like that found in spin glasses, which were discussed in Section 1.2.2.

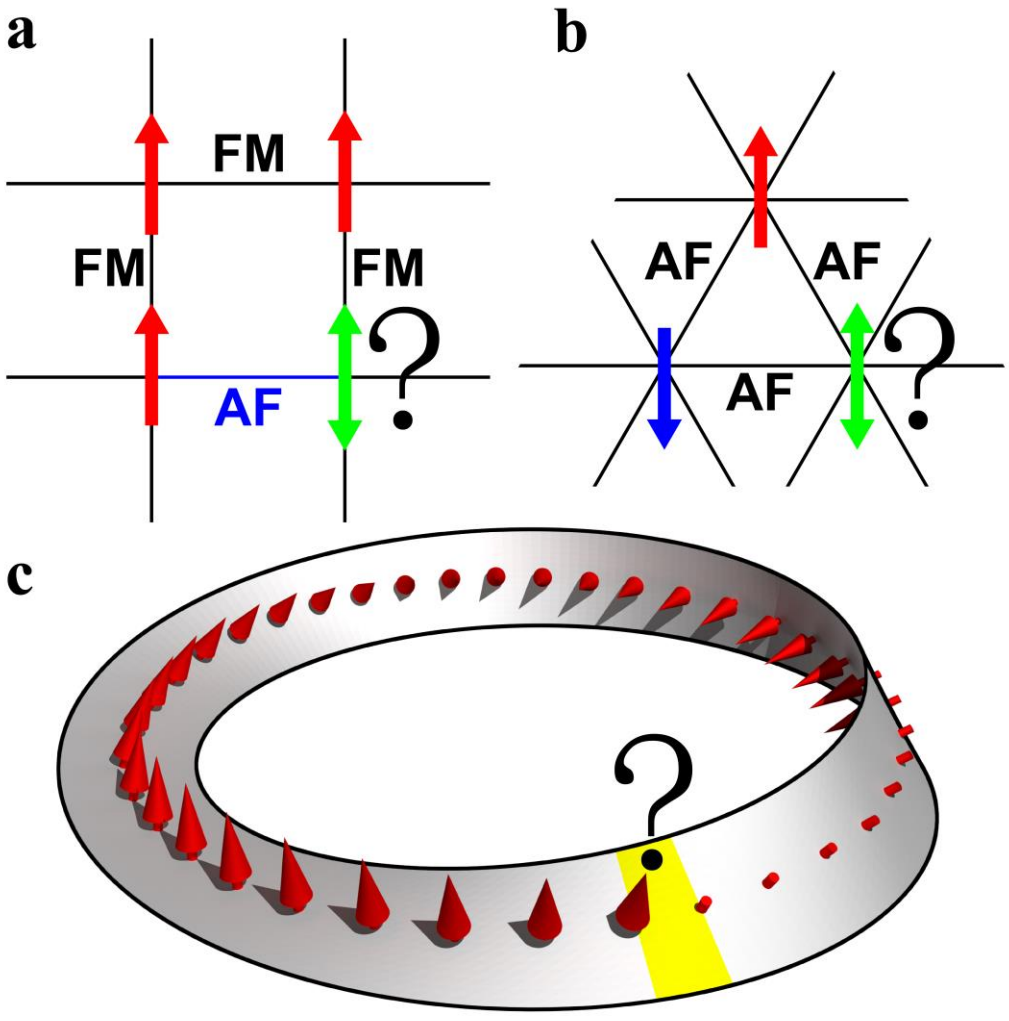

Figure 4.3 **Topological frustration. a-c.** Schematics of different origins of frustration. **a**. Disorder frustration: Disordered bonds between Ising spins on a square lattice cannot be satisfied with any spin configuration, leading to frustration. **b**. Geometric frustration: Antiferromagnetically coupled Ising spins on a triangular lattice are frustrated due to the geometry of the system. **c**. Topological frustration. Ferromagnetic spins perpendicular to a Mobius strip are locally unfrustrated, but the system is still frustrated due to its topology.

As I discussed in Section 1.3, the pyrochlore spin ices are a class of ceramics that have attracted considerable recent scientific interest due to the geometric frustration [13,27] in the magnetic moments of rare earth atoms on the pyrochlore lattice of corner sharing tetrahedra, analogous to Pauling's description of hydrogen disorder in hexagonal water ice at low temperatures [2]. A finite residual entropy at low temperature, resulting from a macroscopically degenerate ground state, is characteristic of these fully frustrated systems. Despite measurements confirming a residual entropy in spin ice [13] and implying a disordered and degenerate low-temperature state, a unique ground state has nevertheless been predicted, based on long-range dipolar interactions in the system [51], but it has proven difficult to observe experimentally [13,53–56]. Notably, one recent study did observe a partial recovery of entropy, below the Pauling value, for single crystal samples of $Dy_2Ti_2O_7$ thermally equilibrated for very long time scales [57]. These considerations are of course influenced by the crystalline perfection of the materials, the nature and density of defects, and the precise configuration of spins around these defects, details that are challenging to experimentally determine or control. ASI presents a possible pathway to resolving these problems as the precise structure of the lattice may be controlled by design.

Square ASI nanomagnet arrays were originally created to model the six-fold plaquette-degeneracy in pyrochlore spin ice [13,75]. The system has proven to be an excellent vehicle for studying frustration and its effects. Unfortunately, it falls slightly short of a completely analogous model to the pyrochlores due to the inequivalent interaction strengths between the four magnets at each vertex. The variation in the distance between adjacent and opposite spins at each vertex lifts the six-fold degeneracy of the ice rule state and leads to a ground state given by a perfect tiling of type I vertices, which is unique up

to a global reversal of spins. In Chapter 2, we demonstrated perfect ground-state ordering in our square ASI samples (see Figure 2.35). Our observation of large domains of ground-state ordering in the square geometry is robust and highly reproducible, demonstrating the benefits of $FePd_3$ as a material system for ASI studies. In this chapter, I discuss results on samples where I have modified the perfect square geometry by introducing edge dislocations into square ASI lattices.

To define our lithography patterns, square lattices are first modeled as a network of connecting springs. The springs are all have equal spring constants and act so as to keep the lengths between neighboring nodes the same, shown schematically in Figure 4.4(a). We also include a term that aims to keep the angles between connections coming out of each node the same, *e.g.* 90° for four connections or 120° for three connections. This is modeled as a torsional spring attached to each pair of connections, as shown in Figure 4.4(b). Equation 4.1 gives a mathematical description of the relaxation algorithm. The force on a specific node *i* is given by

$$\vec{F}_i = \sum_j^{\#NN} k(\vec{x}_i - \vec{x}_j - \vec{x}_0) + \sum_j^{\#NN} \varepsilon(\vec{x}_i - \vec{x}_j) \tag{4.1}$$

where $\vec{x}_n$ are the locations of the different nodes, $\vec{x}_0$ is the equilibrium spring length vector with magnitude of one lattice constant, and k and ε are constants.

To create a defect, we remove a chain of nodes starting from a given point, doubling the spring lengths along the chain of removed points. We then let the system relax, iterating the system by slightly moving each node based on the force acting on it from the other nodes. A damping term (not shown) brings the velocity to zero at the end of each iteration, so that mass and inertial effects may be ignored. Once the lattice has equilibrated, we record the *x* and *y* location of all nodes. We then use these relative coordinates as a guide for our

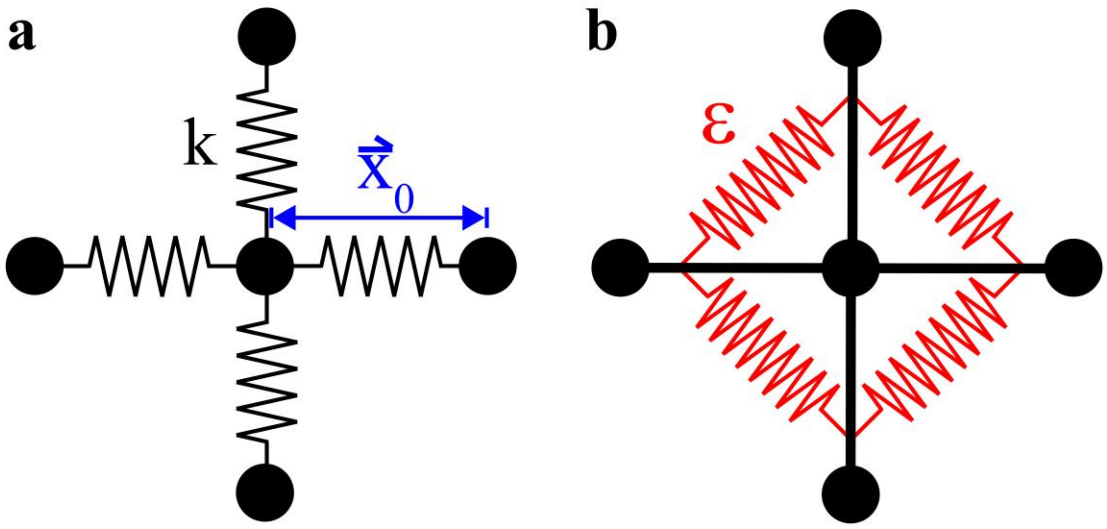

Figure 4.4 **Spring model schematic. a**. Each node is connected to its nearest neighbors by a linear spring of spring constant k and equilibrium length $\vec{x}_0$. **b**. The connections between nodes have torsional springs of spring constant ε attached to them.

lithography patterns. Figure 4.5 shows snapshots of a relaxation process for a lattice with one dislocation. John Cumings developed the relaxation algorithm in MATLAB and wrote the original script to define the x and y locations of the nodes for lattices with one topological defect. I then added to this script the ability to output a lithography pattern file where the element lengths and lattice constant could be controlled. Thomas Marsh, a high school student who worked with me over the summer of 2015, then added the capability to include multiple dislocations in our lattices.

The lattice starts with nodes that are all one lattice constant apart. After the system has fully relaxed, lithography line lengths and locations are found by subtracting a constant amount from the location of the nodes. The line lengths we use here are very similar to the traditional square geometry samples. For the samples in this chapter, the element lengths directly adjacent to the dislocation points needed to be slightly shorter than the other

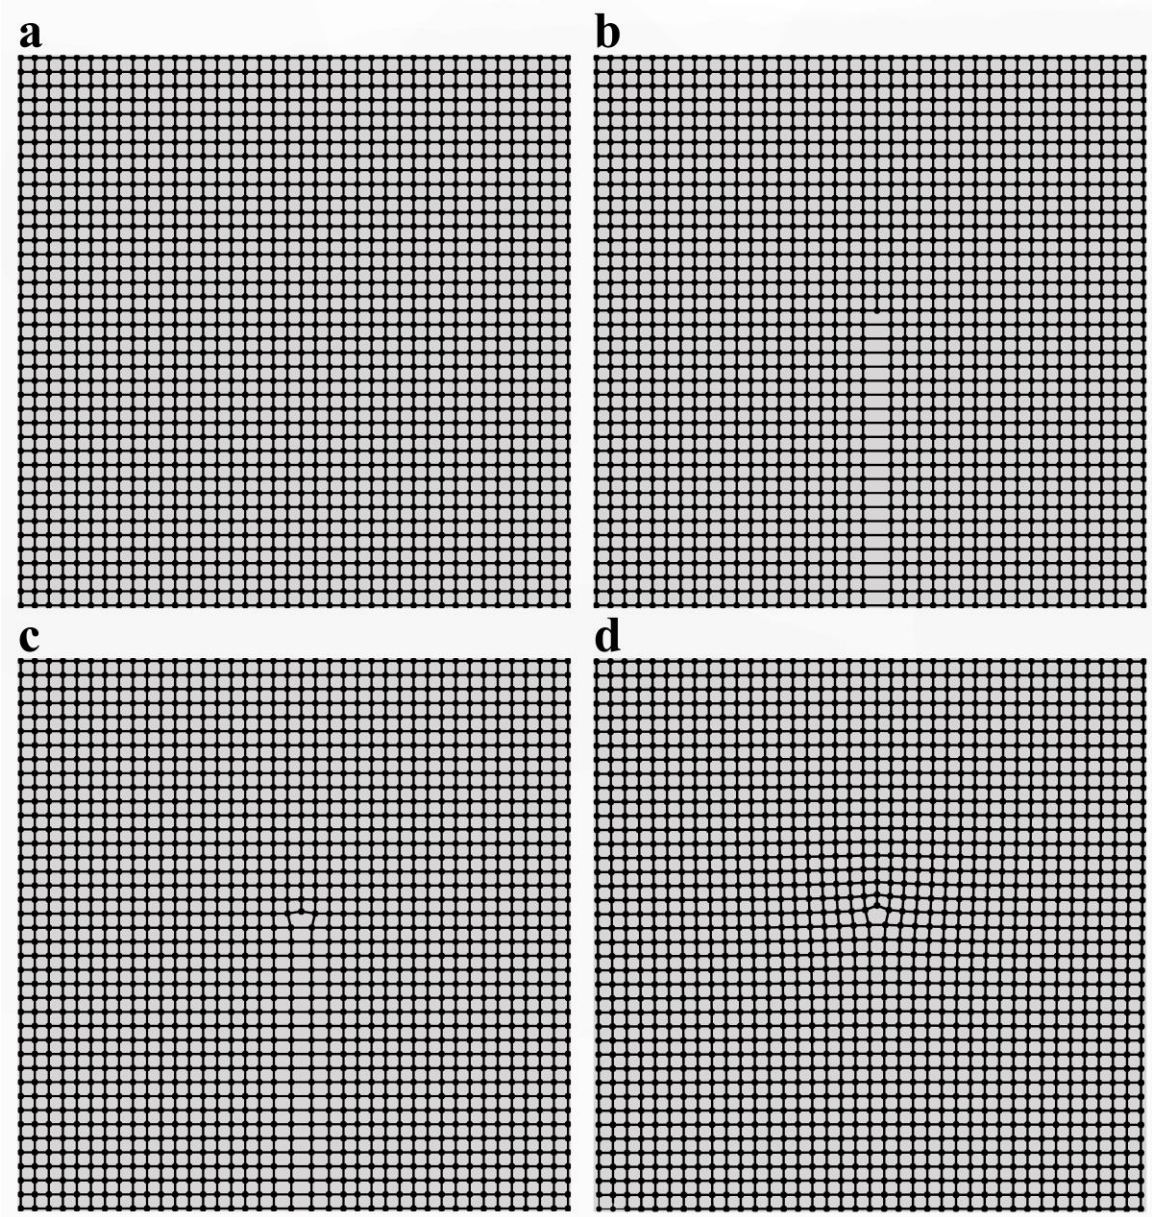

Figure 4.5 **Lithography pattern generation. a-d**. Frames from a relaxation algorithm that generates the *x* and *y* coordinates of the vertices in the lithography patterns for topological defect samples. First, the lattice starts as a perfect square lattice with nodes connected by springs all of equal spring constant (**a**). A chain of atoms is removed starting from a given point (**b**) and the system is then allowed to relax. Frame (**c**) shows the node locations after a few iterations and (**d**) shows the final pattern that is used to generate the locations of the individual spins in each crystal. The springs attempt to keep all the bond lengths the same and our model also tries to keep the angles between the bonds coming out of each node equal as well. See Equation 4.1 in the text for a mathematical description of the model.

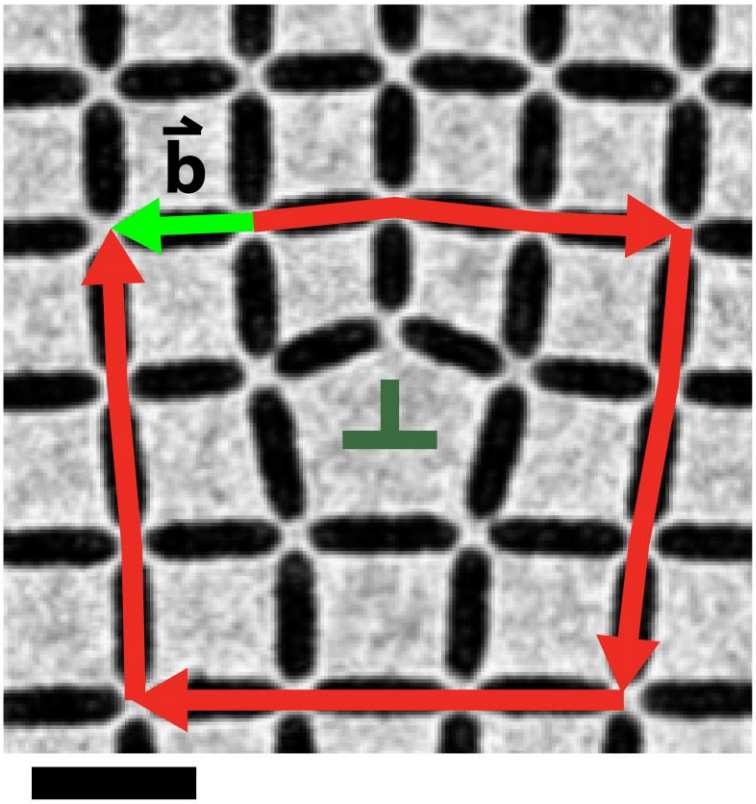

Figure 4.6 **Patterned topological defect**. In-focus TEM image of a section of a square ASI crystal containing a dislocation point defect. The Burgers circuit and vector are overlaid in the image confirming the topological nature of the defect. Scale bar is 500 nm.

elements because the angles around these vertices are much sharper. Additional dose tests were required in order to prevent these elements from touching. Figure 4.6 presents an in-focus TEM image of a dislocation point in our samples. A Burgers circuit is drawn schematically around the defect site to confirm its topological nature.

Figure 4.7 shows a completed full ASI crystal with one topological defect. I fabricated samples with both one and two dislocations in different positions and different distances apart. For the one-dislocation geometries, I will label them with the number 1 to indicate they have one defect and a letter, either *M*, *C*, or *E* to indicate the dislocation is located in the "Middle," "Corner", or "Edge" of the lattice. Within these geometries, the

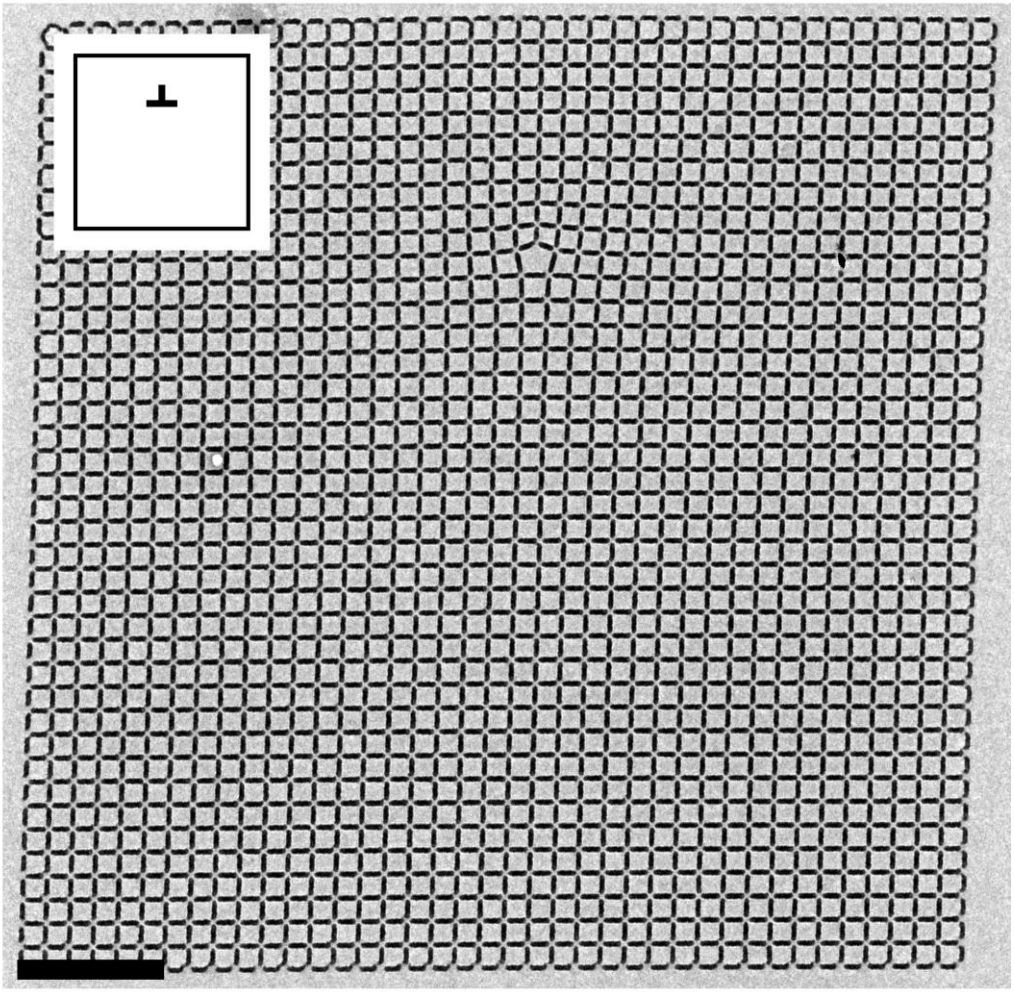

Figure 4.7 **Full crystal with one topological defect**. In-focus TEM image of a full crystal showing one topological defect near the edge, referred to here as a 1-Edge, or 1*E* crystal. Inset shows the schematic representation of this crystal. Scale bar is 3 μm.

dislocation can also point in different crystallographic directions, though we observe no noticeable differences in behavior for the different directions.

I also fabricated many geometries with two dislocations to study how the defects interact and affect each other. Figure 4.8 shows an example of a full crystal with two defects. The two dislocation geometries tended to fall into three main families where, within each family, the distance between the dislocation points was systematically varied. There were some additional patterns studied as well. The three main families are *T* or

"Towards," *A* or "Away," and *R* or "Right." *T* geometries have two defects pointing at each other along the same crystal line in the center of the lattice. These geometries are effectively adding an extra lattice constant along the line between the two dislocations. *A* geometries have two defects pointing away from each other along the same crystal line. This is like removing one lattice constant along the line connecting the defects. *R* geometries have the two defects at right angles to each other. In-focus TEM images of all the crystals belonging to the different families are found in Figures 4.9-4.11. In addition to a letter defining the overall geometry, each crystal type has a number indicating the distance between the two dislocations, expressed in units of a lattice constant, nominally 500 nm.

There are a few other geometries we studied that fall outside of the *T*, *A*, and *R* families. One is shown in Figure 4.8, an *OT* or "Opposite-Towards" configuration. The rest are given in Figure 4.12. These include two *OA* or "Opposite-Away" geometries, as well as a *P* or "Parallel" and a *D* or "Diagonal" geometry. Figure 4.12 also shows an *A*1 geometry, which is not actually a topological defect because it does not have an associated Burgers vector. This geometry represents a relaxed vacancy.

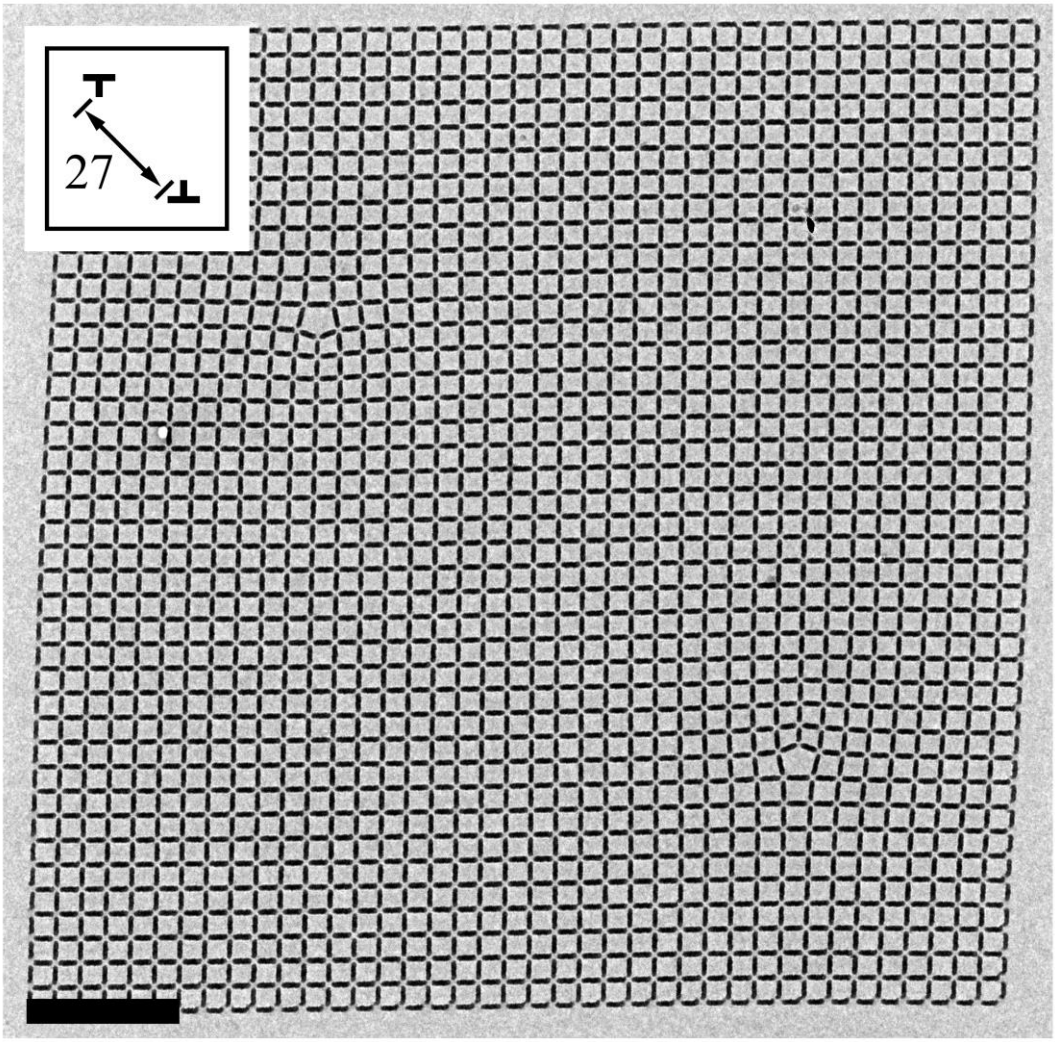

Figure 4.8 **Full crystal with two topological defects**. In-focus TEM image of a full crystal showing two topological defects in an Opposite-Towards, or *OT* configuration. This crystal has the defects 27 lattice constants apart and thus the geometry is defined as *OT*27. The inset shows a schematic representation of the geometry indicating the location of the dislocation points, their directions, and the distance between them. Scale bar is 3 μm.

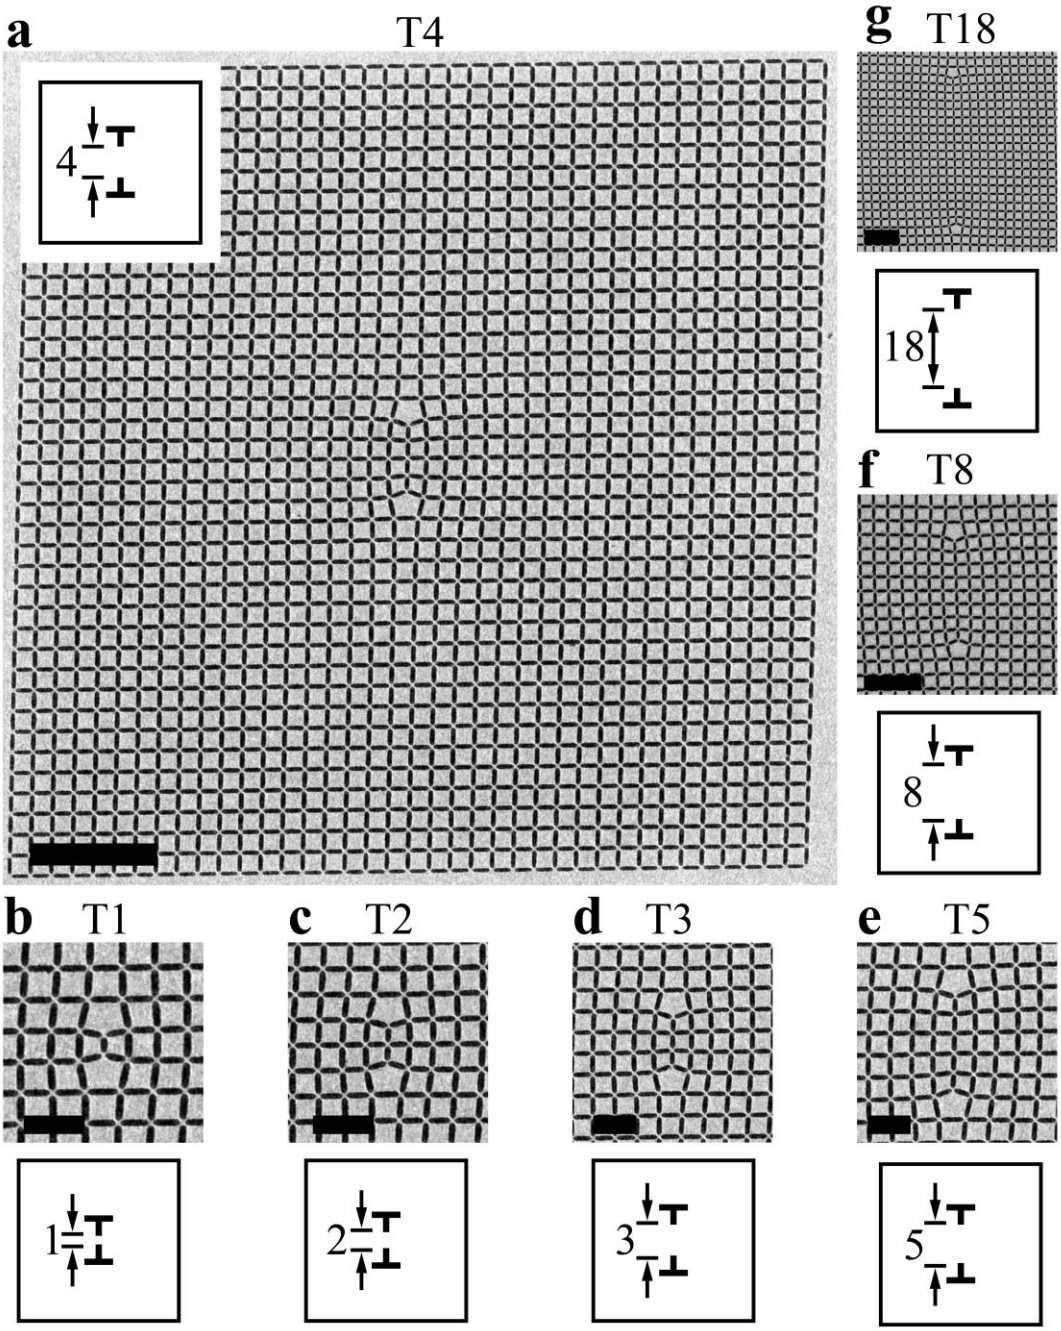

Figure 4.9 **Family of *T* geometries**. In-focus TEM images of all of the "Towards", or *T*, geometries studied. **a**. Full *T*4 crystal with 4 lattice constants between the defect points where the dislocations are pointing towards each other along the same crystal line. Inset is the schematic representation of the crystal showing the location of the dislocation points, their directions, and the distance between them. Scale bar is 3 μm. The other frames show a section of the crystal around the defects and the associated schematic for that geometry. The other geometries studied are (**b**) *T*1, (**c**) *T*2, (**d**) *T*3, (**e**) *T*5, (**f**) *T*8, and (**g**) *T*18. Scale bars for (**b**)-(**e**) are 1 μm, scale bars for (**f**) and (**g**) are 2 μm.

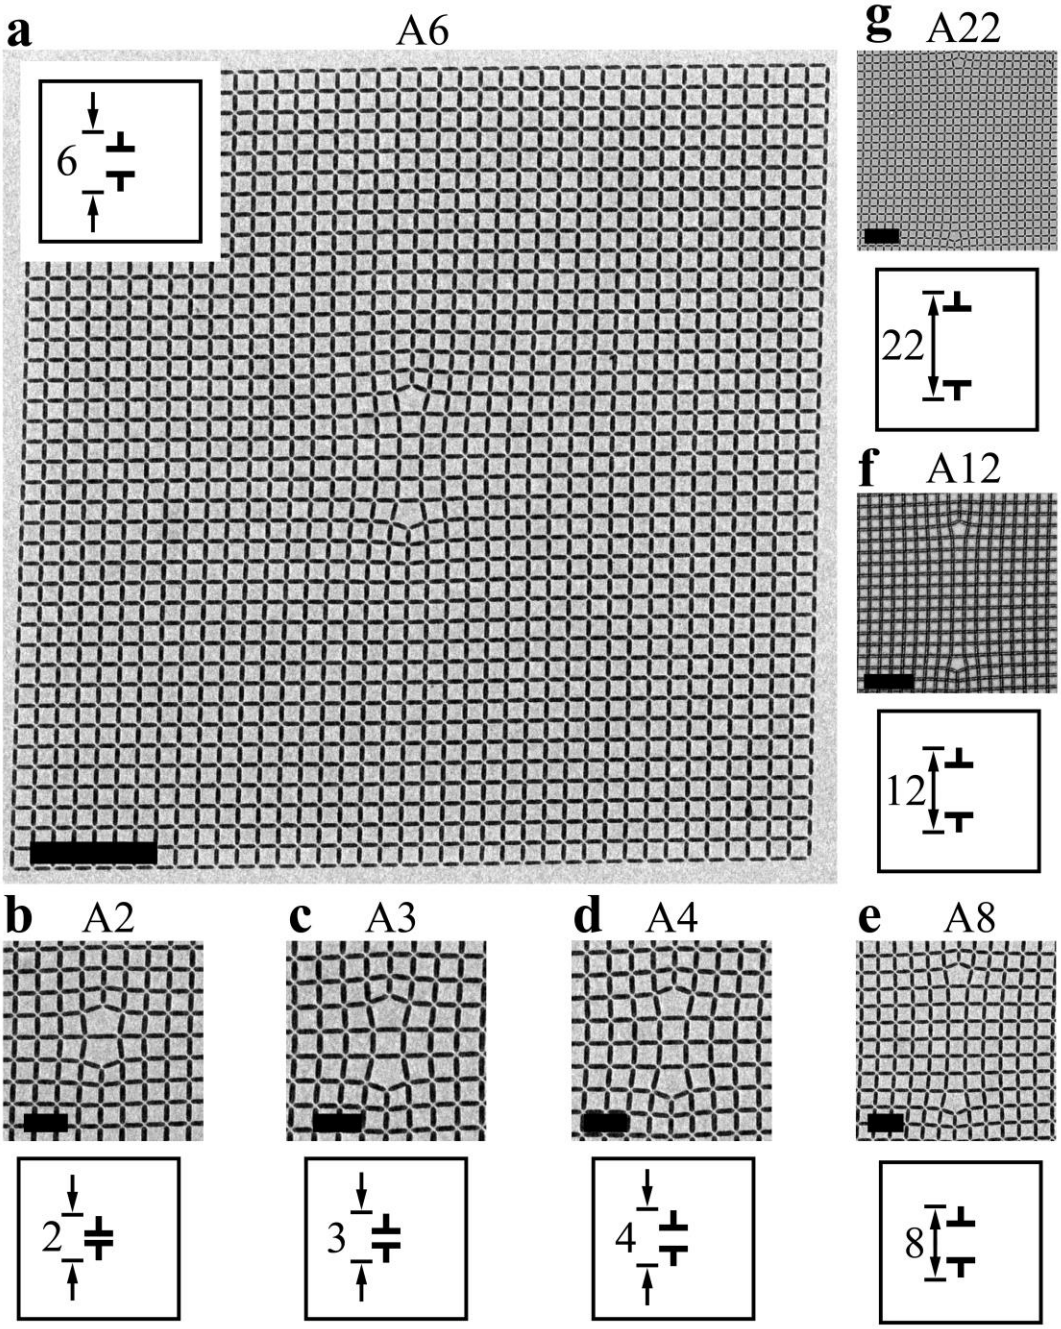

Figure 4.10 **Family of *A* geometries**. In-focus TEM images of all of the "Away", or *A*, geometries studied. **a**. Full *A*6 crystal with 6 lattice constants between the defect points where the dislocations are pointing away from each other along the same crystal line. Inset is the schematic representation of the crystal showing the location of the dislocation points, their directions, and the distance between them. Scale bar is 3 μm. The other frames show a section of the crystal around the defects and the associated schematic for that geometry. The other geometries studied are (**b**) *A*2, (**c**) *A*3, (**d**) *A*4, (**e**) *A*8, (**f**) *A*12, and (**g**) *A*22. Scale bars for (**b**)-(**e**) are 1 μm, scale bars for (**f**) and (**g**) are 2 μm. Note that (**f**) is actually a Lorentz TEM image as there were no high quality in-focus images of the *A*12 geometry readily available.

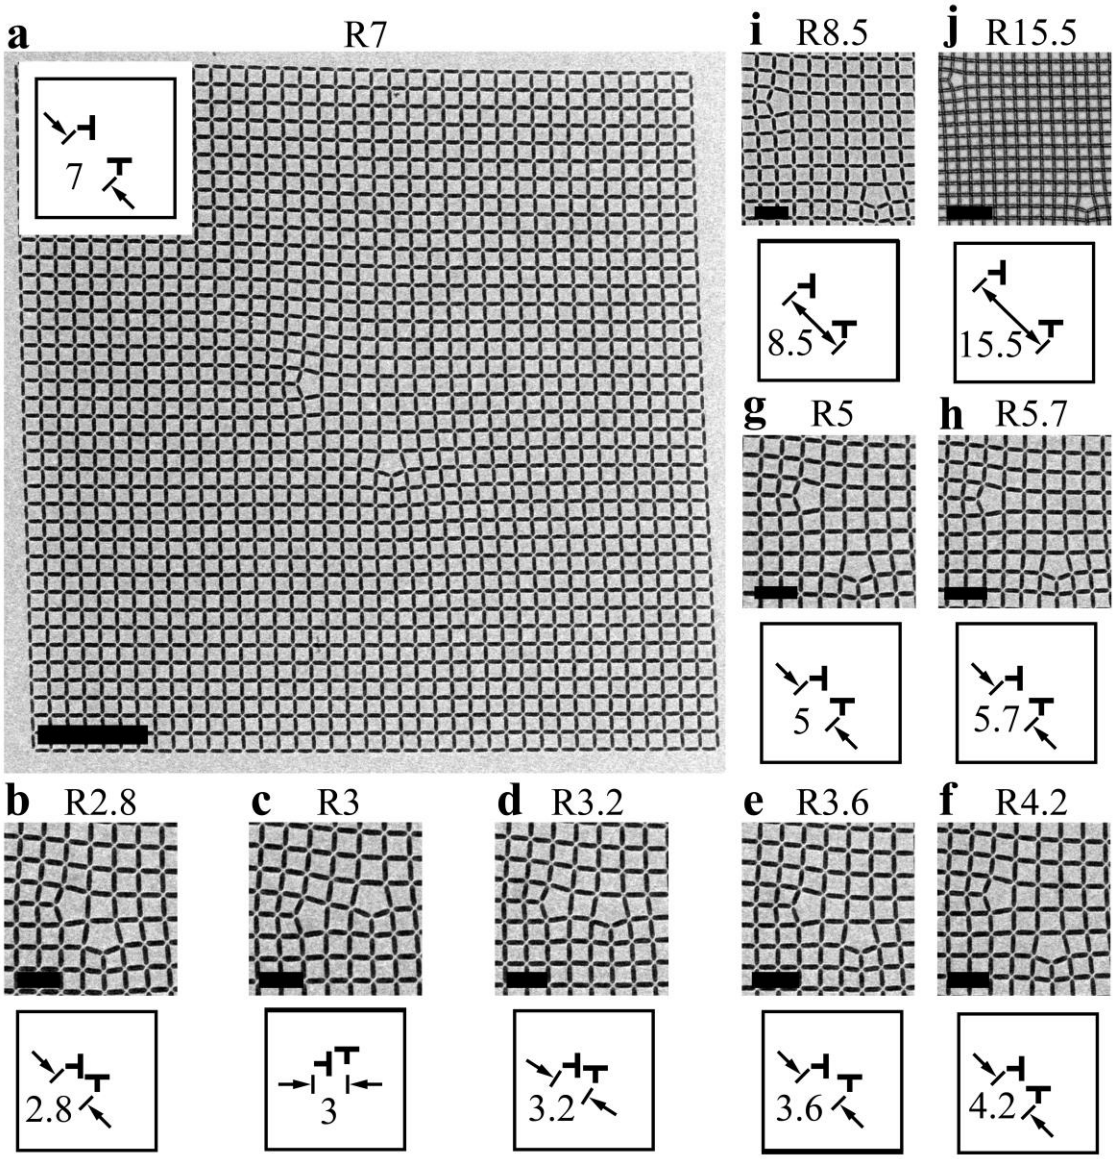

Figure 4.11 **Family of *R* geometries**. In-focus TEM images of all of the "Right", or *R*, geometries studied. **a**. Full *R*7 crystal with 7 lattice constants between the defect points where the dislocations are at right angles to each other. Inset is the schematic representation of the crystal showing the location of the dislocation points, their directions, and the distance between them. Scale bar is 3 μm. The other frames show a section of the crystal around the defects and the associated schematic for that geometry. The other geometries studied are (**b**) *R*2.8, (**c**) *R*3, (**d**) *R*3.2, (**e**) *R*3.6, (**f**) *R*4.2, (**g**) *R*5, (**h**) *R*5.7, (**i**) *R*8.5, and (**j**) *R*15.5. Scale bars for (**b**)-(**i**) are 1 μm, scale bar for (**j**) is 2 μm. Note that (**j**) is actually a Lorentz TEM image as there were no high quality in-focus images of the *R*15.5 geometry readily available.

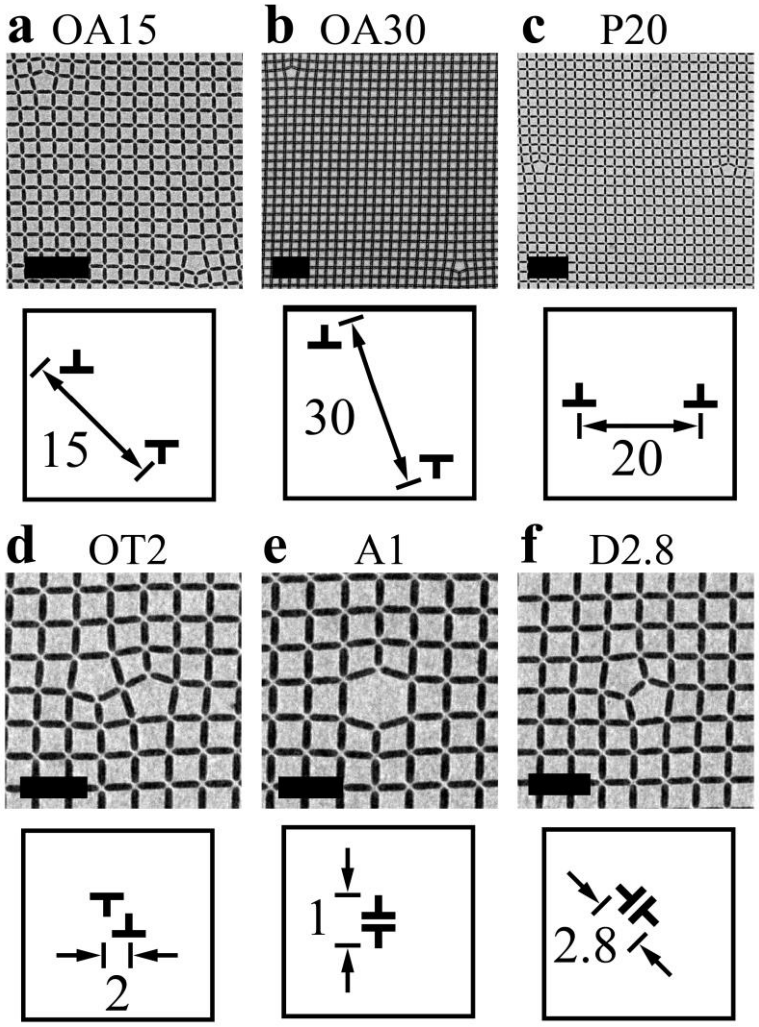

Figure 4.12 **Other assorted geometries studied**. In-focus TEM images of all of other geometries studied. Each frame shows a section of the crystal around the defects and the associated schematic for that geometry. The other geometries studied are (**a**) *OA*15, (**b**) *OA*30, (**c**) *P*20, (**d**) *OT*2, (**e**) *A*1, and (**f**) *D*2.8. Scale bars for (**a**)-(**c**) are 2 μm, scale bars for (**d**)-(**f**) are 1 μm. Note that (**b**) is actually a Lorentz TEM image as there were no high quality in-focus images of the *OA*30 geometry readily available.

Similar to the thermal ASI experiments discussed in Chapters 2 and 3, samples discussed in this chapter were heated using the annealing protocol in Figure 2.33 and Lorentz TEM images were captured after degaussing the microscope's objective lens. The MATLAB crawler for analyzing the Lorentz TEM images was modified by Thomas Marsh to work for all of the dislocation geometries. This was extremely helpful as the script was

able to automatically handle any of the geometries presented to it without additional input. The vertex notation shown in Figure 2.36 and used extensively throughout this chapter was also implemented by Marsh. I will also note here that some of the TEM images in this chapter, both in-focus and Lorentz contrast, have a skew distortion in the image. This is because all of the images are taken in the low magnification mode of the microscope. In this mode, the focal length of the electron optics is longer, and since image distortions scale with the focal length of the lens, the distortion is pronounced in many of my images [182]. The skew is fairly small and can be corrected with digital image processing software, though this process is tedious and does not add any new information. Image distortion can be seen, for example, in Figures 4.11(a) and 4.19, as well as many other images in this chapter. Corrected images can be found, for example, in Figures 4.14, 4.17, and 4.45.

In this chapter, we essentially introduce a new kind of frustration, specifically due to the topology of the system. We find the presence of a topological defect results in extended frustration within our samples. Locally, the magnets are unfrustrated, but frustration of the lattice persists due to its topology and strictly prevents the perfect square ground-state ordering found in Figure 2.35. ASI was originally developed to model real frustrated systems and while many interesting results and insights have been gained from studying ASI, it has rarely been able to make predictions about and aid in understanding real crystalline systems. In this chapter, we aim to accomplish this and use our samples to speculate on the impact of topological defects in a variety of systems. After discussing the behavior of our ASI lattices at the single spin level, we argue that topological defects could explain some unsolved problems in real systems and deserve more attention in future work on a range of materials with non-zero order parameters.

4.2 Results and discussion

4.2.1 Main results

In this section, we present extensive results and observations on topological frustration of ASI, providing the first realization of domain wall pinning due to a lattice defect with information at the single-spin level. We discuss our observations in the context of crystalline materials systems with non-zero order parameters and develop a theoretical model for the presence of topological frustration within a system. Our main results are summarized in Figure 4.13. We find that crystals with dislocations show large domains of ground-state ordering, but they also always have a chain of type II and type III higher energy vertices originating from the dislocation point. These chains are required because of the topological nature of the dislocation. A closed path around the defect is topologically altered, making it impossible to support continuous ground-state order, and thus a frustrated chain must be present. The chains in the single-dislocation geometries always continue to an edge of the finite crystal, as shown by representative examples in Figures 4.14-4.16 for the 1*E*, 1*M*, and 1*C* geometries. Figure 4.16 shows a nice example of a domain wall originating from the defect site and not taking the shortest path to the edge of the crystal, a behavior we observe in many different crystals. The 1*C* geometry has a dislocation very close to the edge, only 8 lattice constants away, but the vertex chain extends away from the closest edge. Figure 4.15 shows a very well ordered crystal with minimal defects and domains. This crystal is mainly split into two domains with an additional small one in the bottom right corner. Figure 4.14 displays a few domains in

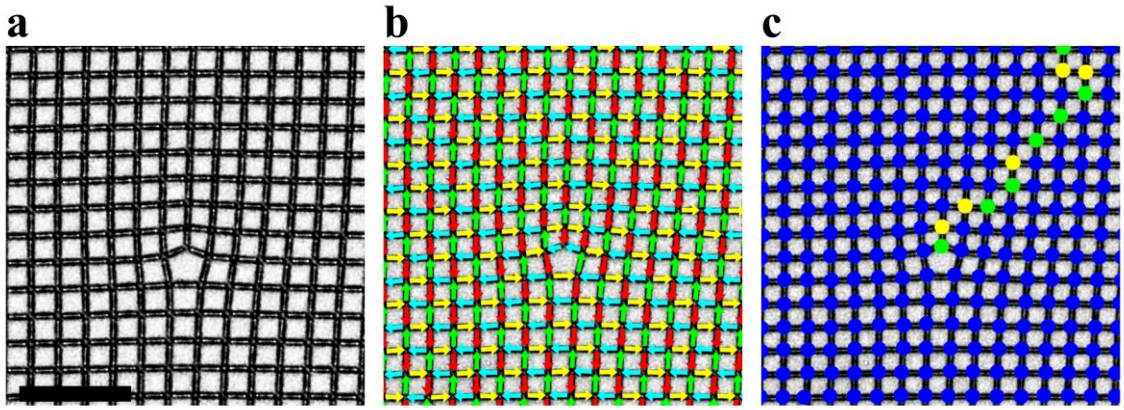

Figure 4.13 **Topological defect after annealing. a**. Lorentz contrast TEM image of the same defect site as in Figure 4.6 after annealing. **b**. The same Lorentz contrast TEM image in (**a**) with arrows overlaid indicating the direction of each macro spin. The image shows mainly ground-state order with a domain wall originating from the dislocation point. **c.** Vertex map of the same image in (**a**) and (**b**) representing the same information, but with vertex types as defined in Figure 2.36. Two and three spin vertices found in our samples are simply mapped onto the corresponding four-spin type, excluding the highest energy, type IV vertex because we do not observe any in our samples. This notation clearly shows a chain of type II and type III vertices originating from the defect site surrounded by a single domain of ground-state order. Scale bar is 2 μm.

addition to the topologically required vertex chain. These crystals are representative of the behavior we see in one-dislocation geometries.

The crystals shown in Figures 4.14-4.16 display the extended frustration of the system through long, meandering vertex chains within large regions of ground-state order. Other domains and domain walls may be present as are seen in our crystals of canonical square ASI, for example, in Figures 2.37 and 2.38, and in other previous studies as well [78], but a vertex chain must always begin at the dislocation point. Other than this, the frustrated chains and conventional domain walls are indistinguishable, exhibiting the same basic phenomena. The lowest energy configuration for either would be the shortest possible, straight chain of only type II vertices, extending in a <11> direction, but this is not generally observed. Type III vertices are almost always present in both the frustrated

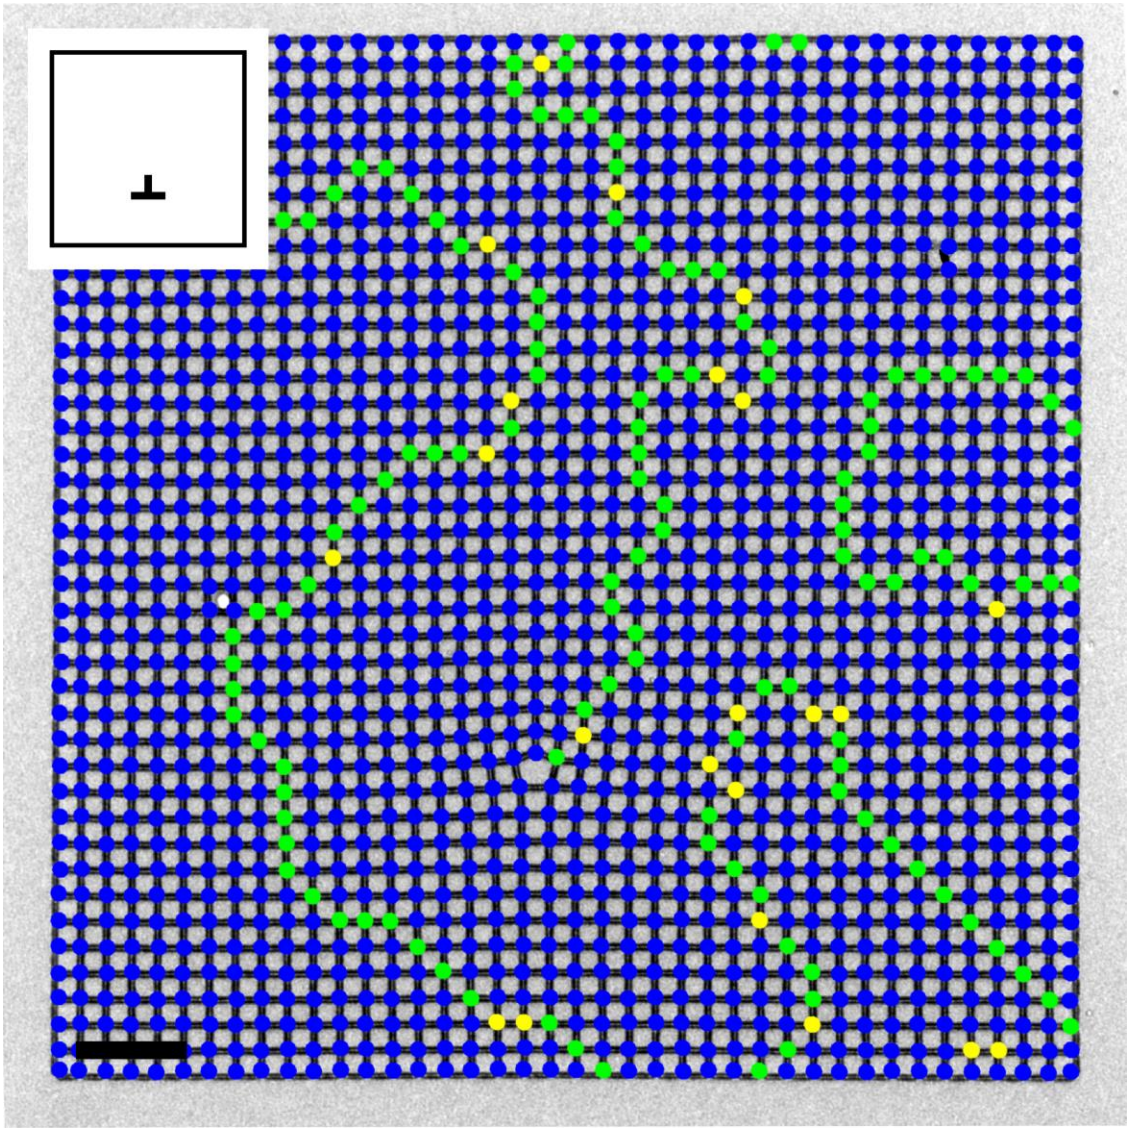

Figure 4.14 **Vertex map of a crystal with one topological defect.** Lorentz contrast TEM image with vertex types overlaid on a 1*E* geometry crystal. A meandering vertex chain originates from the dislocation among a few ground-state domains separated by traditional domain walls. The inset shows the schematic representation of this geometry indicating the location and direction of the dislocation. Scale bar is 2 μm.

chains and the conventional domain walls likely because the type III vertices play an essential role in domain wall motion in the square ice lattice [119]. Furthermore, since a frustrated wall's presence is required by topology, there is only a fractionally small change in its total energy to include type III vertices and take a longer, often bent or winding path

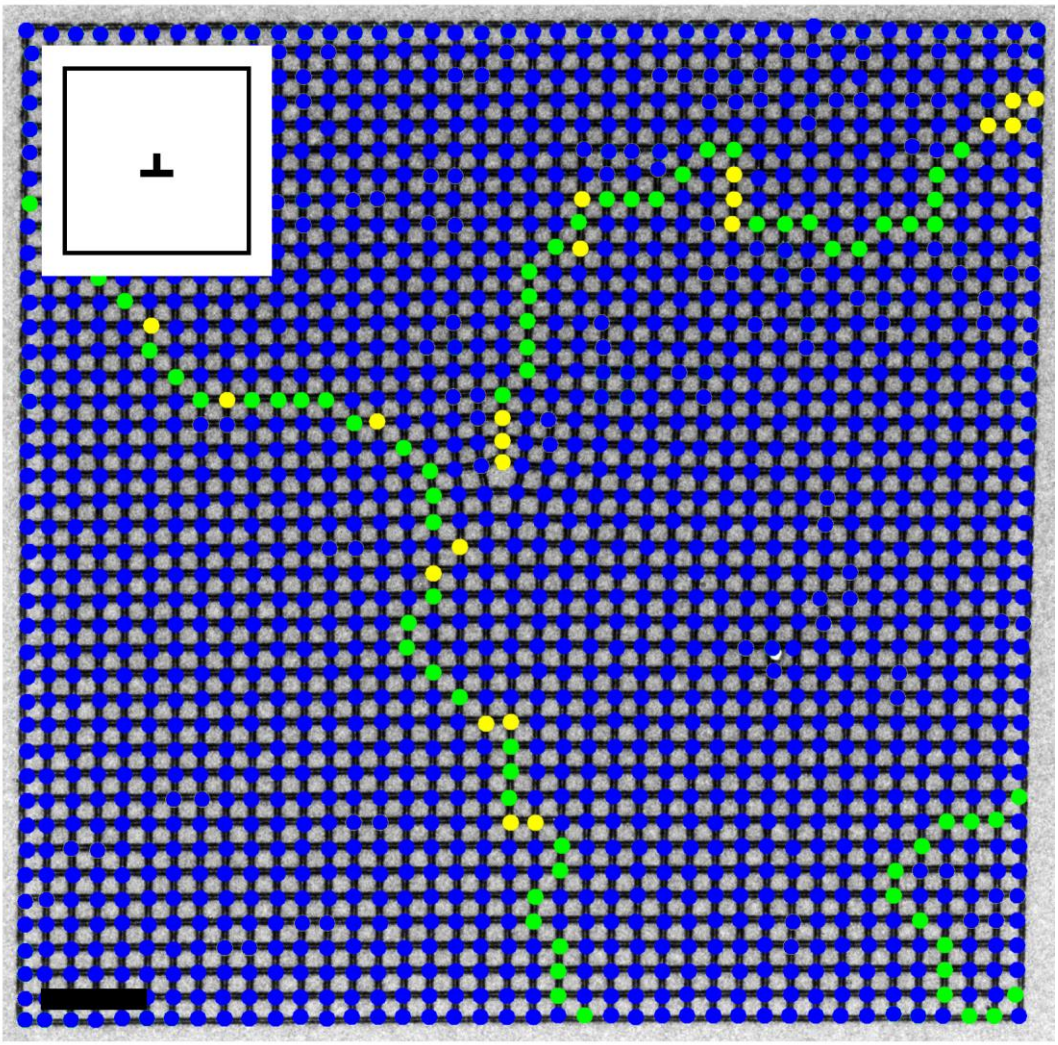

Figure 4.15 **Vertex map of a crystal with one topological defect.** Lorentz contrast TEM image with vertex types overlaid on a 1*M* geometry crystal. A meandering vertex chain originates from the dislocation among a few ground-state domains separated by traditional domain walls. The inset shows the schematic representation of this geometry indicating the location and direction of the dislocation. Scale bar is 2 μm.

through the lattice. It appears that both such walls are at least partly entropically driven, as opposed to purely energetically driven, as the lowest energy configurations do not dominate and the walls in general do not appear to have any tension. Additionally, we note here that topological frustration is distinct from vertex frustration, which is found, for

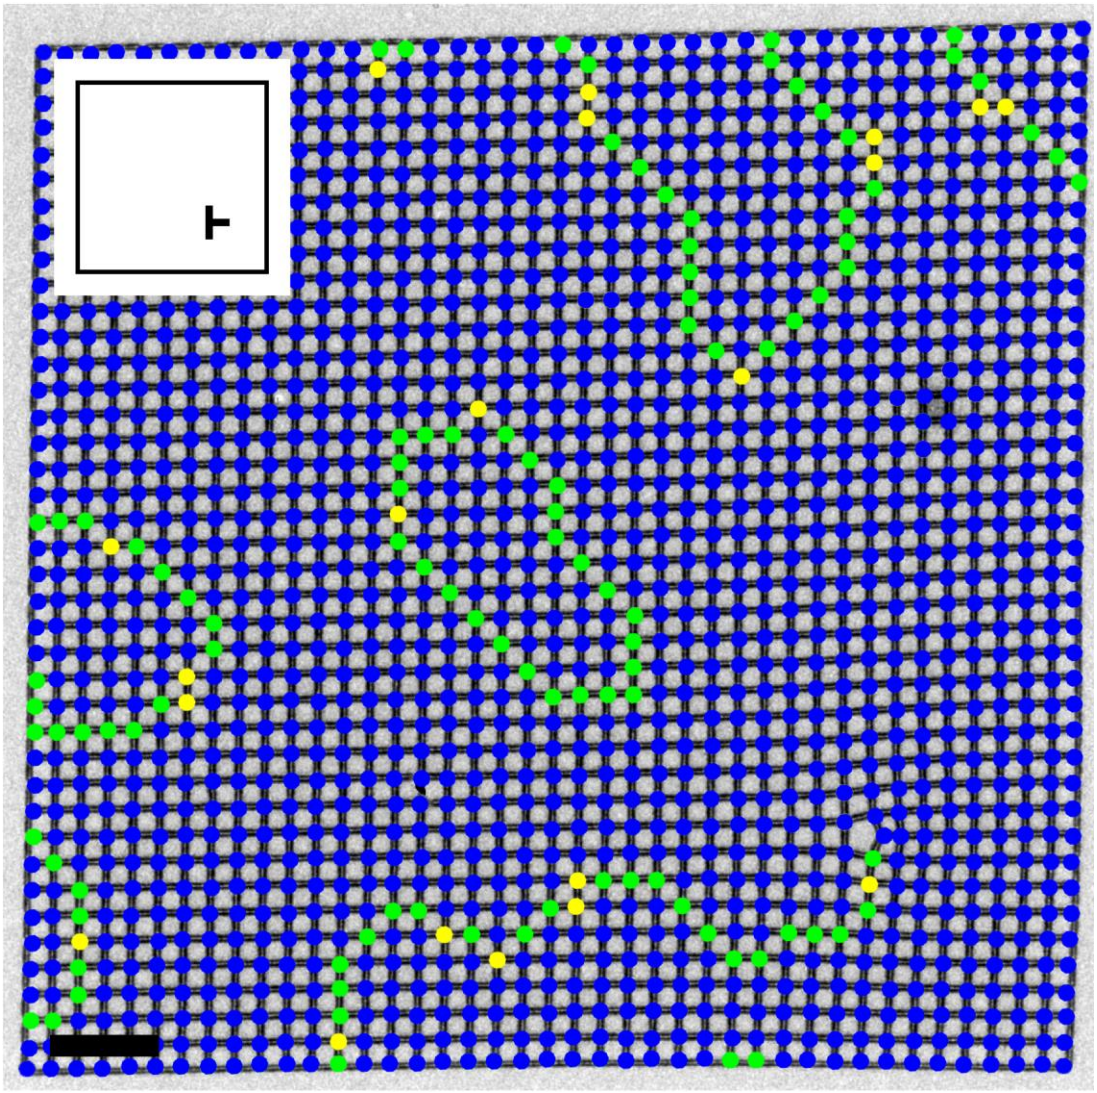

Figure 4.16 **Vertex map of a crystal with one topological defect.** Lorentz contrast TEM image with vertex types overlaid on a 1*C* geometry crystal. A meandering vertex chain originates from the dislocation among a few ground-state domains separated by traditional domain walls. The inset shows the schematic representation of this geometry indicating the location and direction of the dislocation. Scale bar is 2 μm.

example, in the shakti lattice [79] because vertex frustration is only due to the geometry of the system, not its topology.

For crystals with two dislocations, we see a wide range of behaviors. Typically, and especially when the defects are more than a few lattice constants apart, the two defects have separate domain walls that each meander to the edges of the crystal. Occasionally, a

single domain wall will connect the two dislocations, as the inclusion of two topological defects actually allows for the possibility of a larger continuous ground-state domain. Examples of both of these effects are shown in Figure 4.17. In both one- and two-dislocation geometries, the crystal is effectively a regular square geometry crystal far away from the defects, but the topological nature of the dislocations can still have long-range effects that support domain walls permeating the whole lattice.

We have fabricated a variety of two-dislocation geometries with varying distance between the defects as discussed in Section 4.1. When the two dislocations are extremely close to each other, they can essentially be thought of as one defect. It is straightforward to draw a Burgers circuit around both defects and find that no vector is required to complete the loop and naively conclude that there is no topological defect inside. In reality, there are two topological defects inside with equal magnitude but opposite direction Burgers vectors that cancel each other out. Taking a loop in the same direction (clockwise or counterclockwise) around each defect produces two Burgers vectors that add to 0. In these defects that are extremely close together, we still see both domain walls that connect the defects and separate domain walls that individually extend to the edges of the crystal. Examples of some well-ordered crystals that display these alternating behaviors for very close defects are given in Figures 4.18-4.21. It should be surprising that even at essentially no separation, the two dislocations can with some frequency behave as individual defects that appear not to interact or to know about other defects in the lattice, and they independently nucleate domain walls to accommodate the topological frustration in the system.

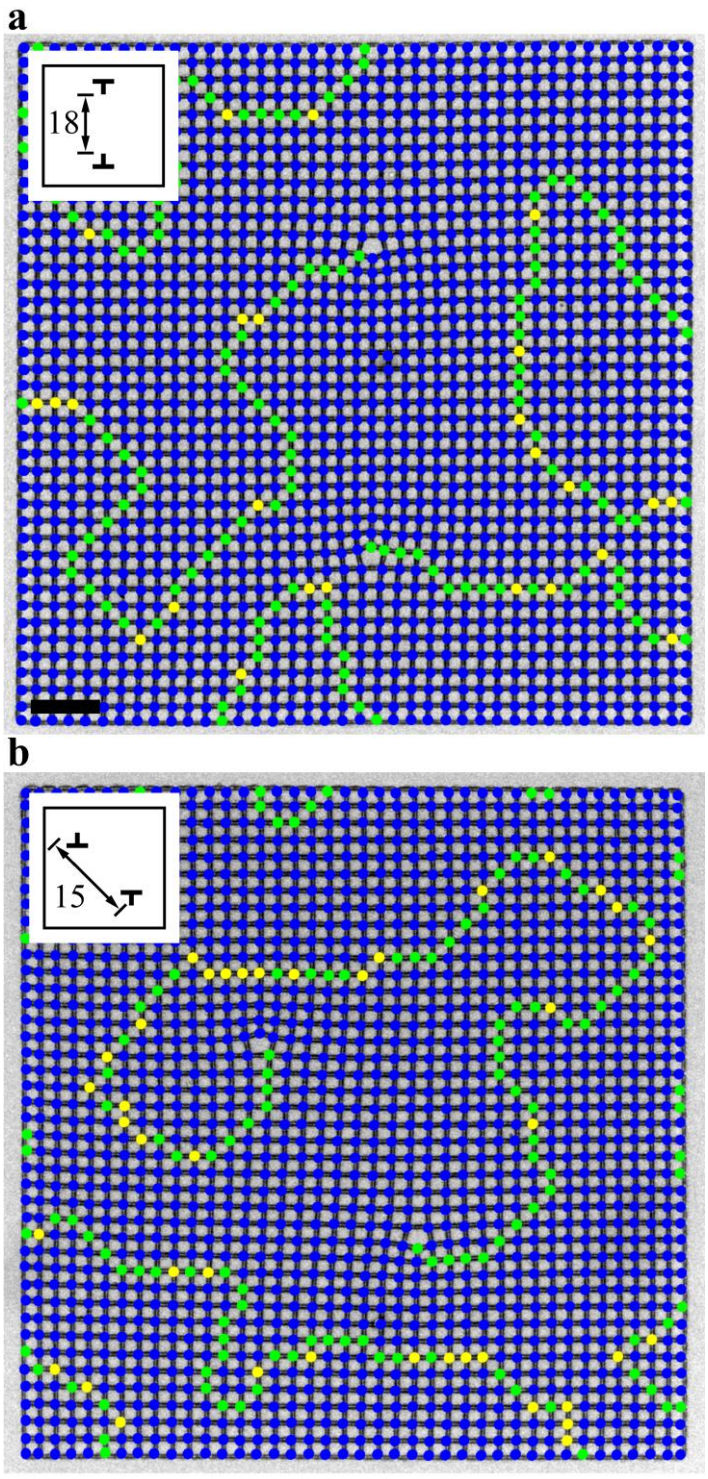

Figure 4.17 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on a (**a**) *T*18 and (**b**) *OA*15 geometry crystal. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Insets show the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

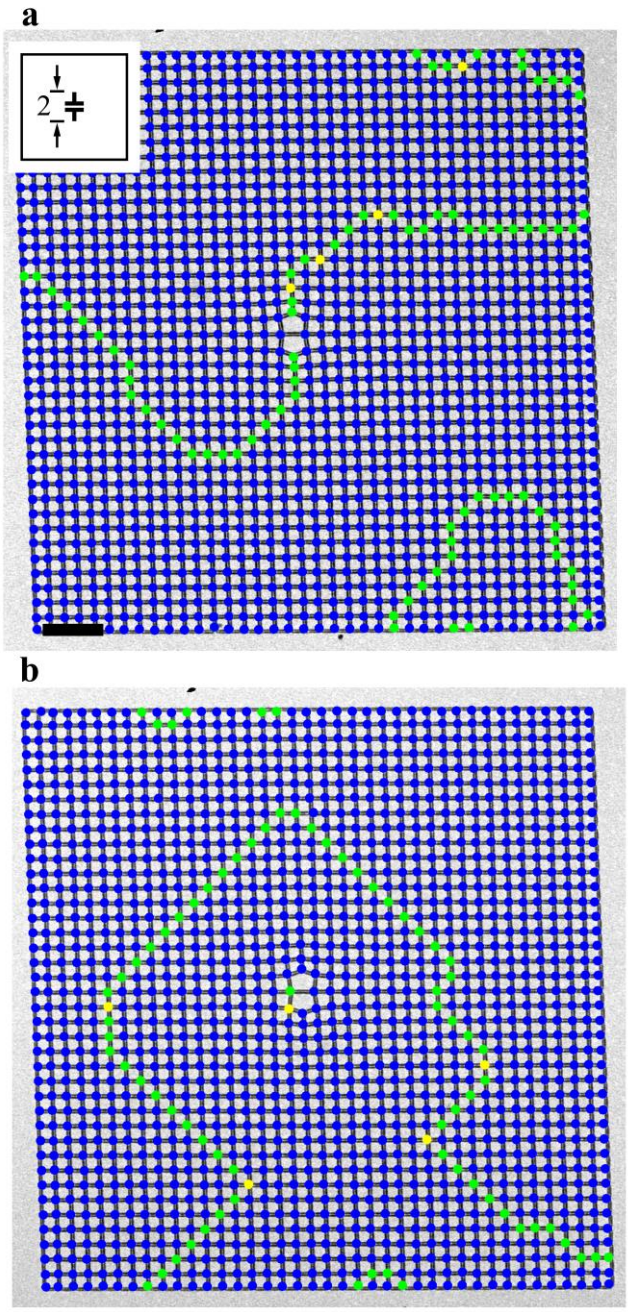

Figure 4.18 **Vertex maps of crystals with two topological defects spaced very close together.** Lorentz contrast TEM images with vertex types overlaid on *A*2 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. The defects are still topological defects because a Burgers circuit and vector can be defined for each one individually, though a loop around both of them does not indicate the presence of any topological anomaly because they cancel each other out. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

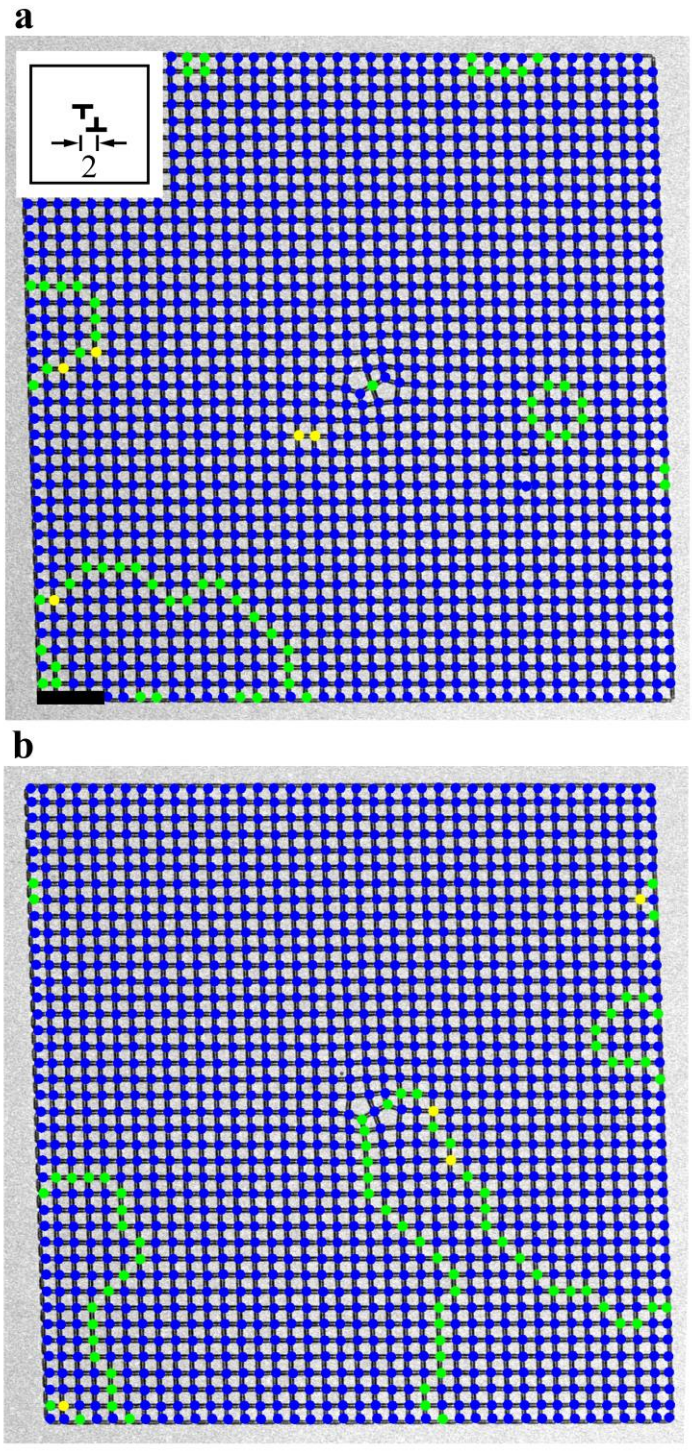

Figure 4.19 **Vertex maps of crystals with two topological defects spaced very close together.** Lorentz contrast TEM images with vertex types overlaid on *OT*2 geometry crystals showing domain walls (**a**) connecting and (**b**) not connecting the dislocations. In (**a)**, a single type II vertex connects the two dislocation. In (**b**), domain walls begin at each dislocation and end at the edges of the crystals. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

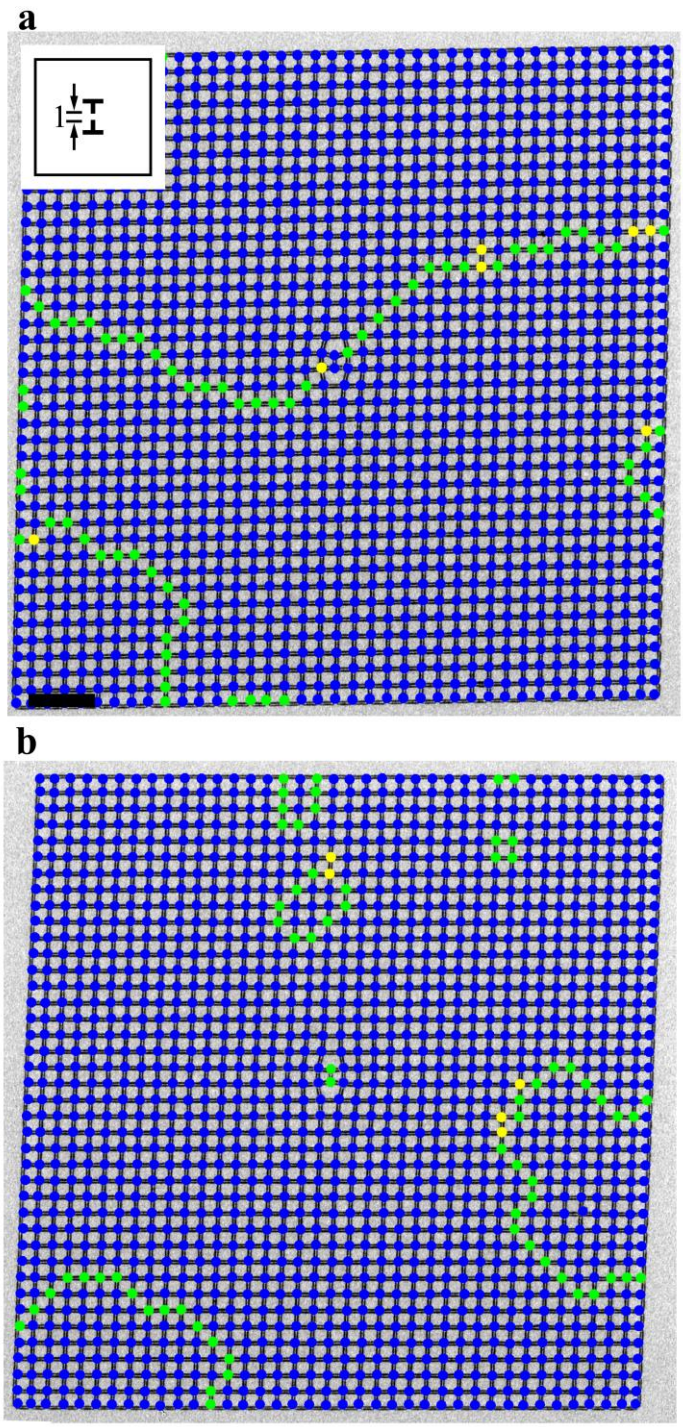

Figure 4.20 **Vertex maps of crystals with two topological defects spaced very close together.** Lorentz contrast TEM images with vertex types overlaid on *T*1 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

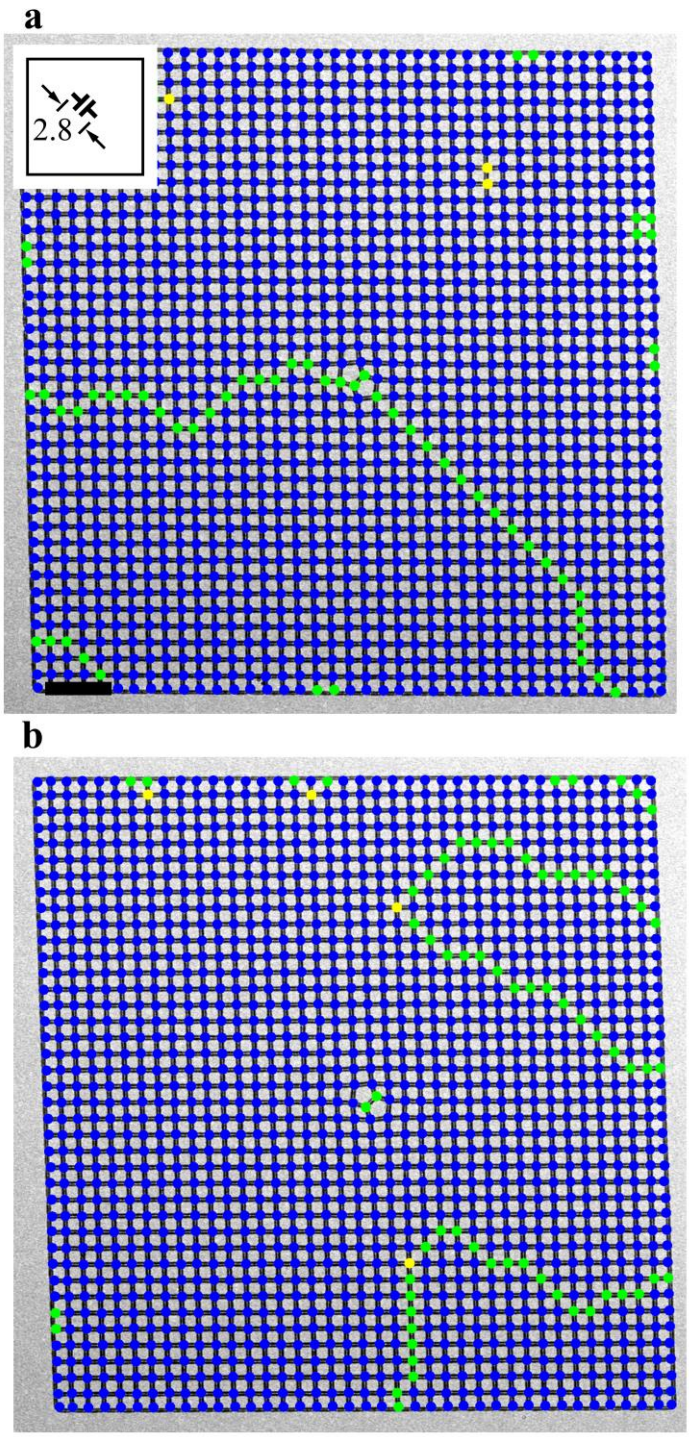

Figure 4.21 **Vertex maps of crystals with two topological defects spaced very close together.** Lorentz contrast TEM images with vertex types overlaid on *D*2.8 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a)**, domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

As we move the defects further away from each other, we still see domain walls that connect the two defects, but this tends to become less and less likely as the dislocations get further apart. Here we present a representative, but not exhaustive, selection of well-ordered crystals that show a range of behaviors. Figures 4.22-4.32 display a variety of two-dislocation geometries, each with an example of domain walls not connecting the dislocations and an example of a crystal where a domain wall does connect the dislocations. There are a lot of images that all show similar effects over many of the different geometries studied. Note that some of the examples here, such as Figures 4.22(a) or 4.26(a), are perfectly ordered except for the topologically required domain wall. Some other crystals, such as in Figures 4.27(a) or 4.28(b), are nearly perfect except for the required wall. We also see that Figure 4.32(b) has an interesting domain wall that snakes in between and around the defects to connect them.

We can conclude from these images that the inclusion of two dislocations creates extensive domain walls in the system, but the exact geometry and placement of the dislocations are less important to the overall behavior, especially when the defects are more than a few lattice constants apart. Behaviors for the different geometries at the closest spacings will be discussed in more detail below. We also note that the domain walls themselves do not interact with each other or form “X” or “Y” junctions, which is an interesting observation in its own right.

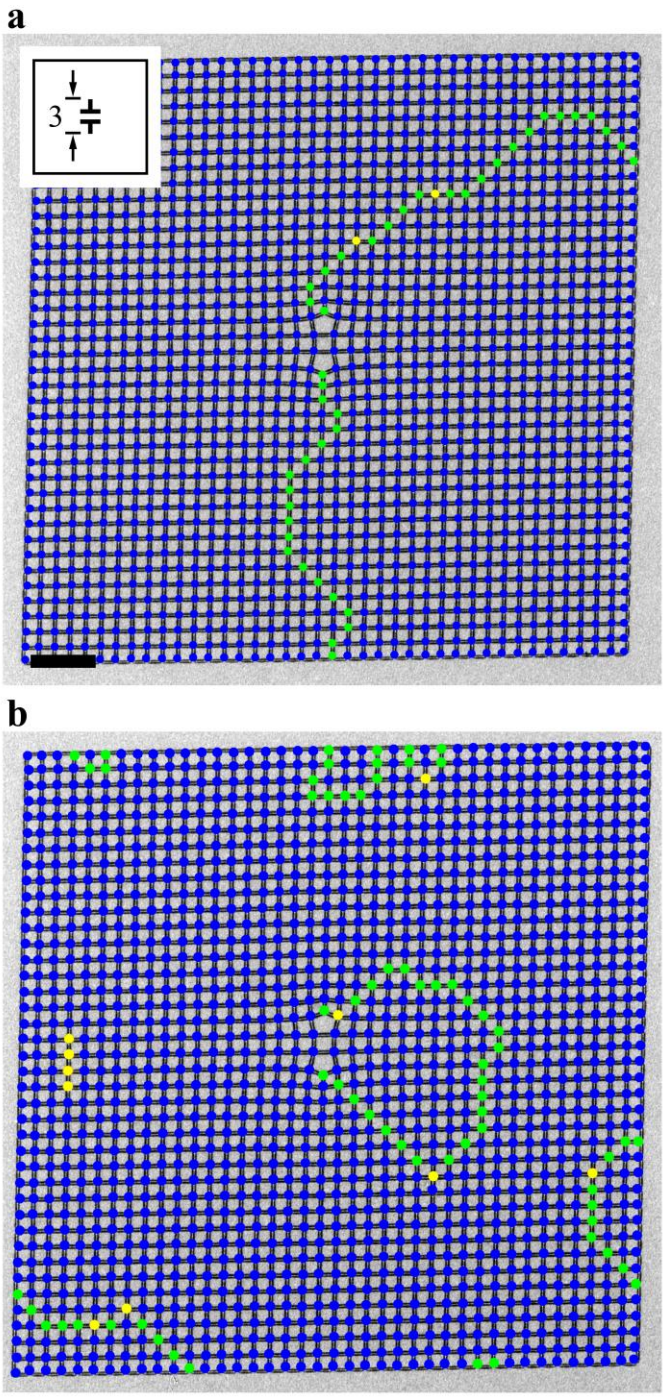

Figure 4.22 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *A*3 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

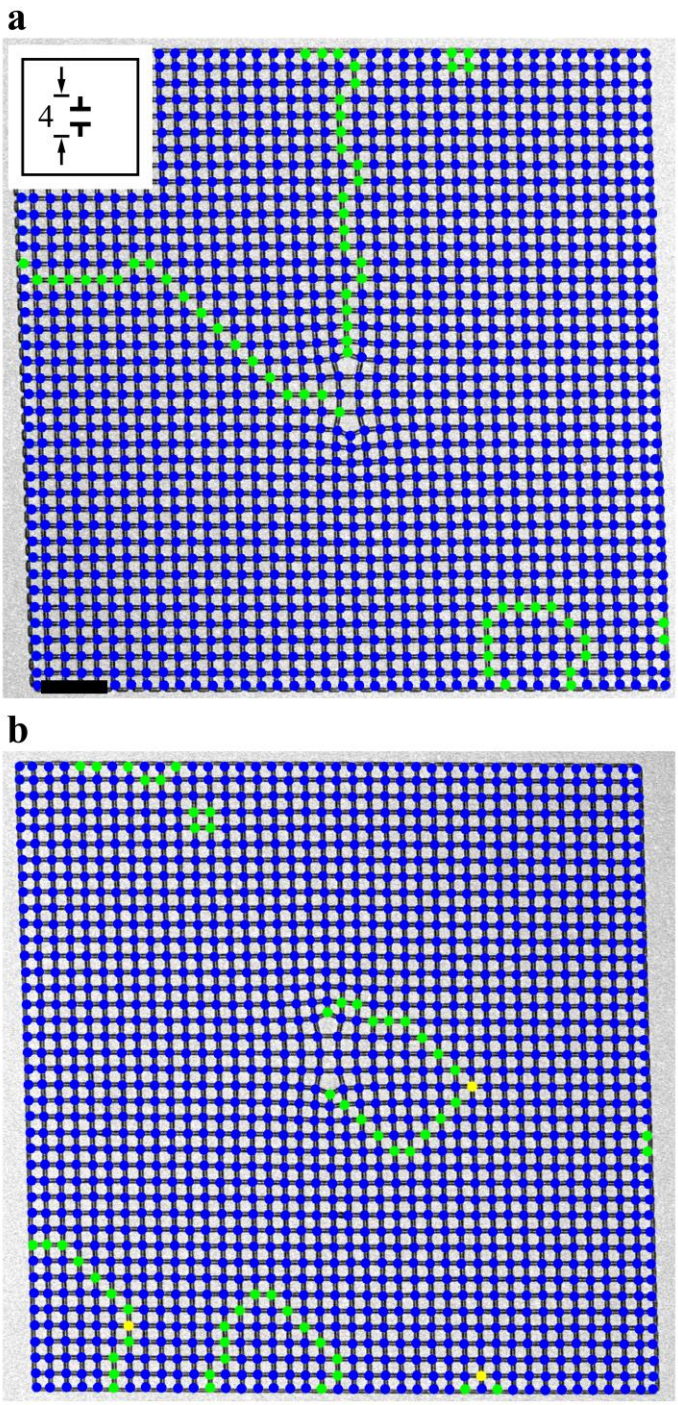

Figure 4.23 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *A*4 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

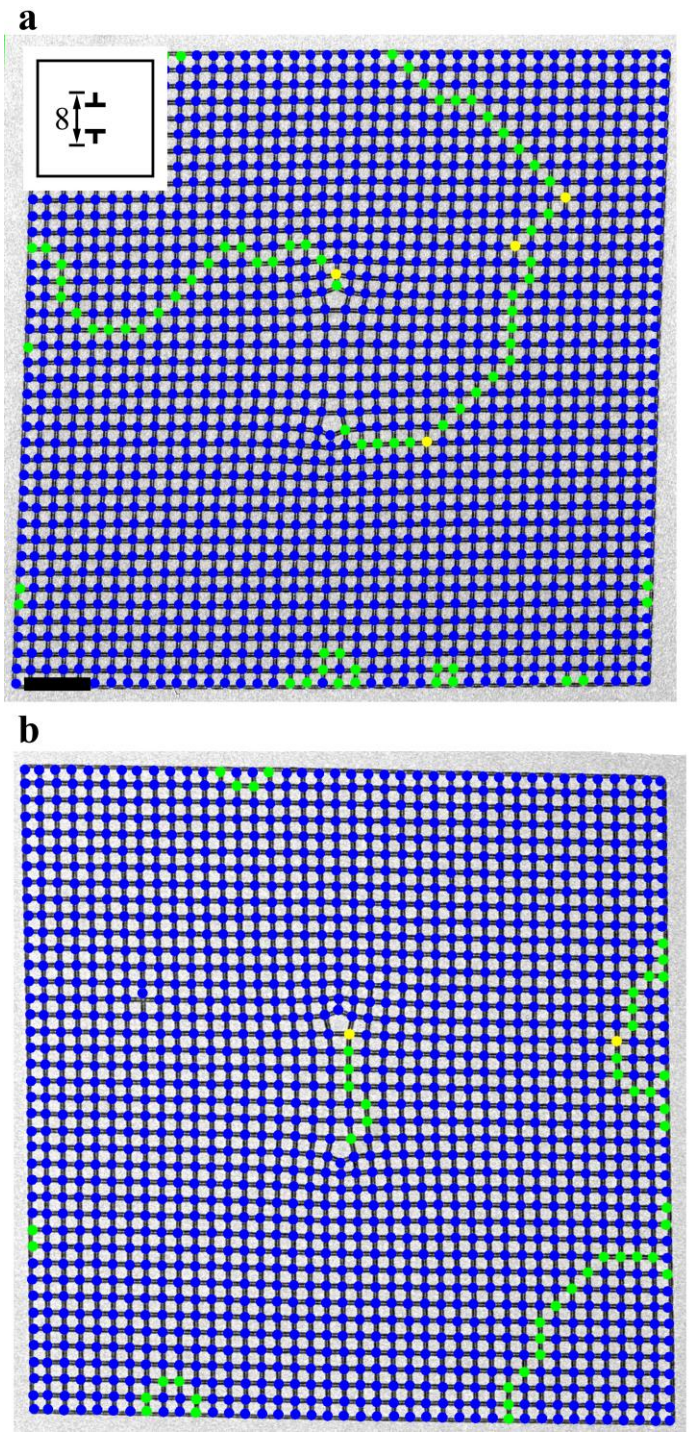

Figure 4.24 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *A*8 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

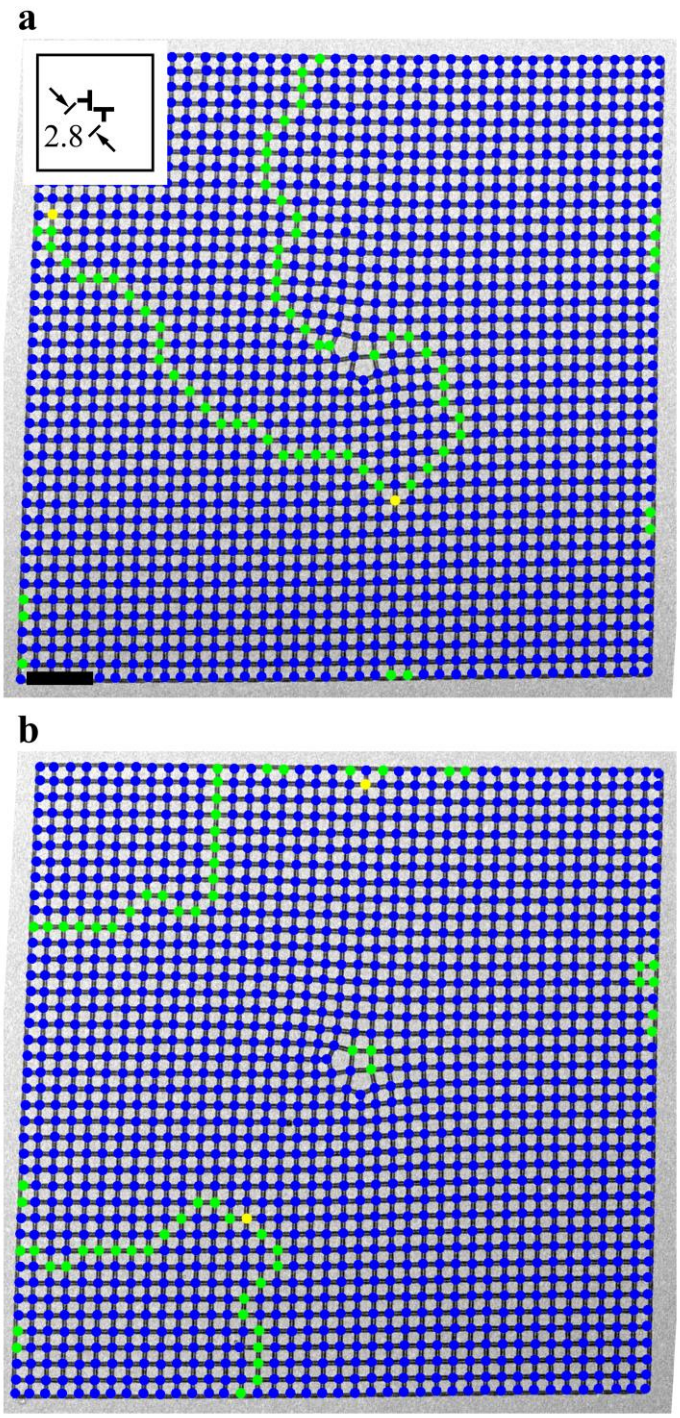

Figure 4.25 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *R*2.8 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

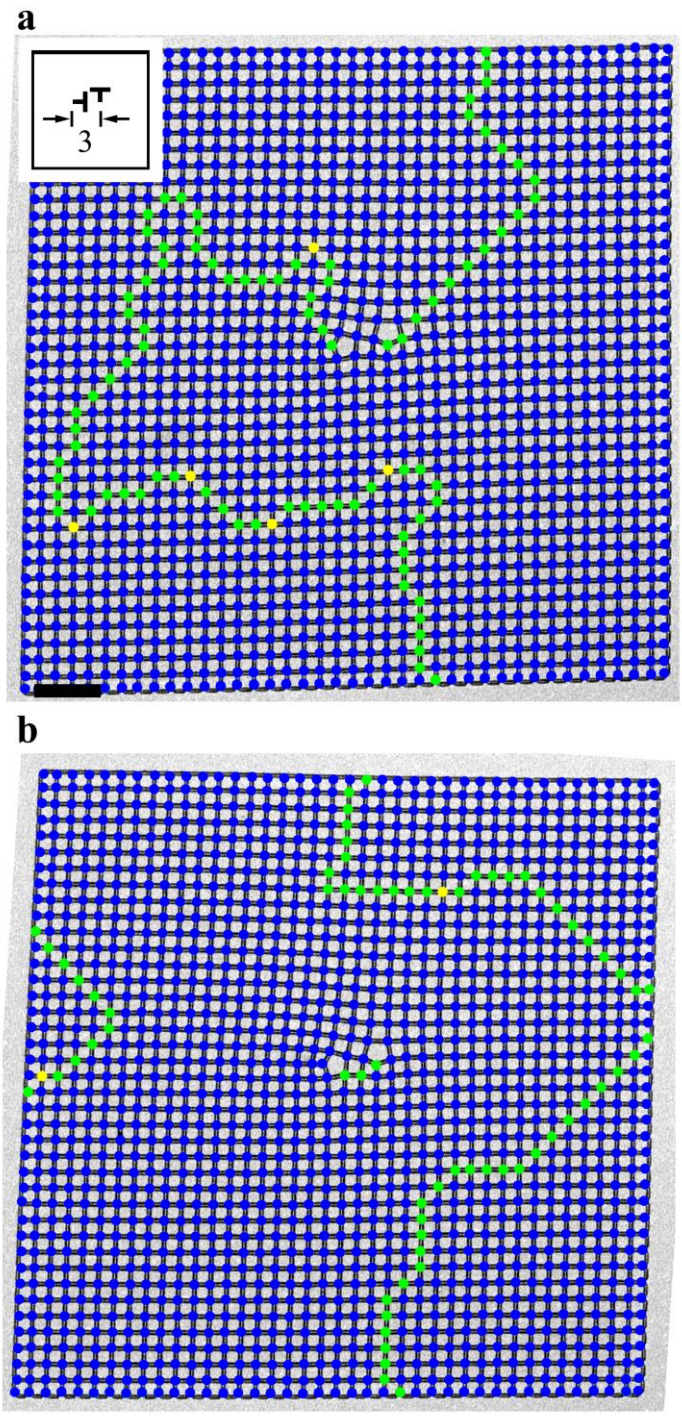

Figure 4.26 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *R*3 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

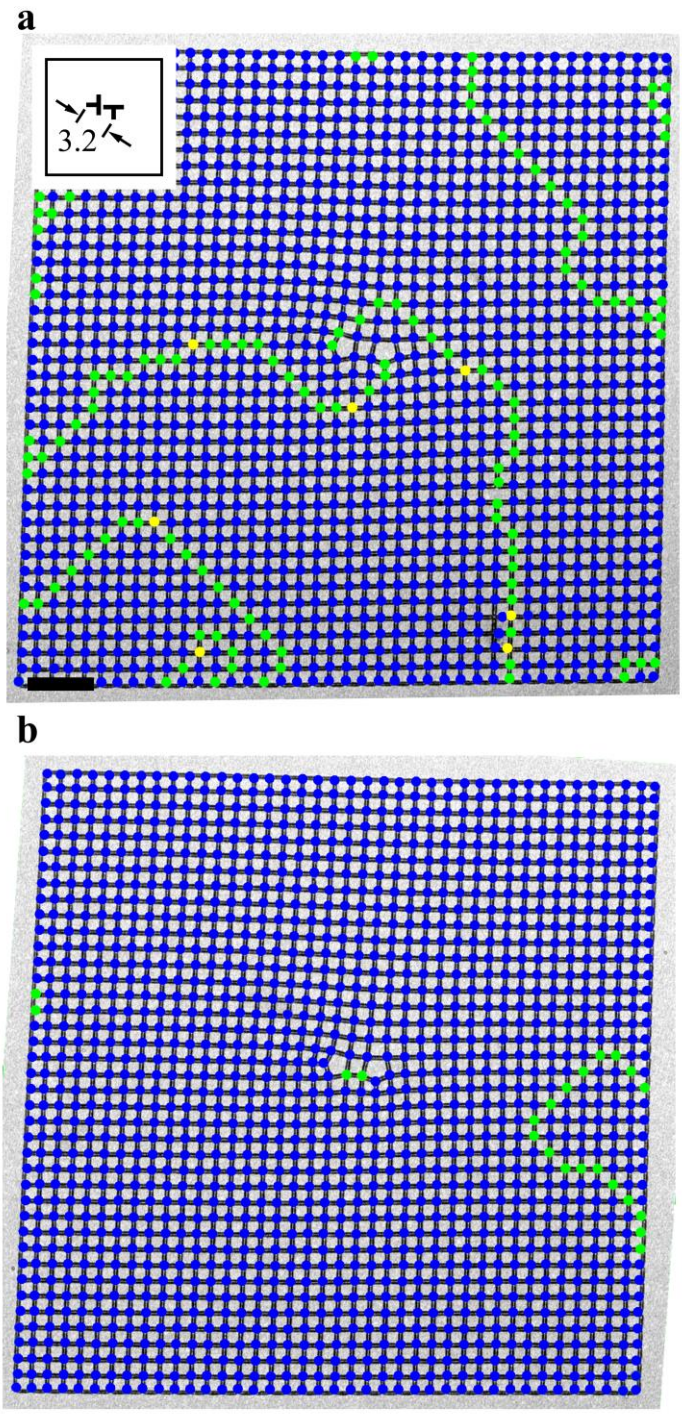

Figure 4.27 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *R*3.2 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

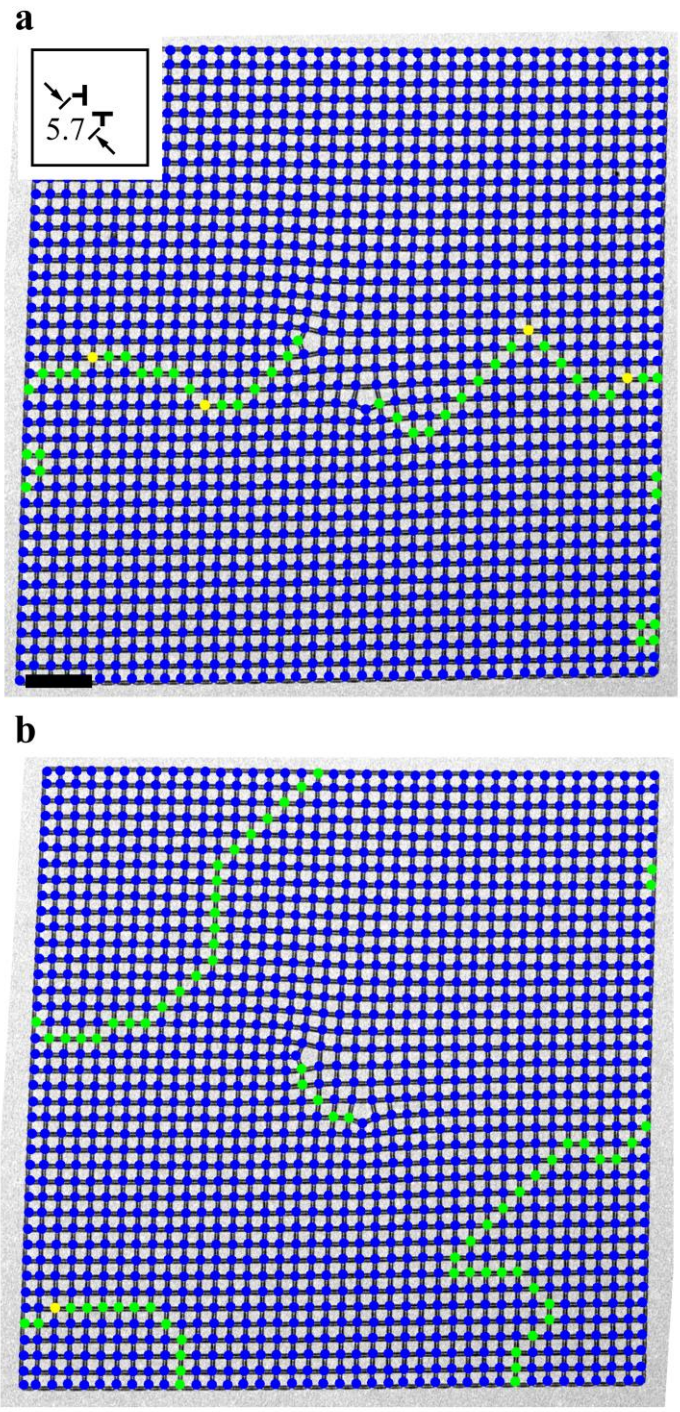

Figure 4.28 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *R*5.7 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

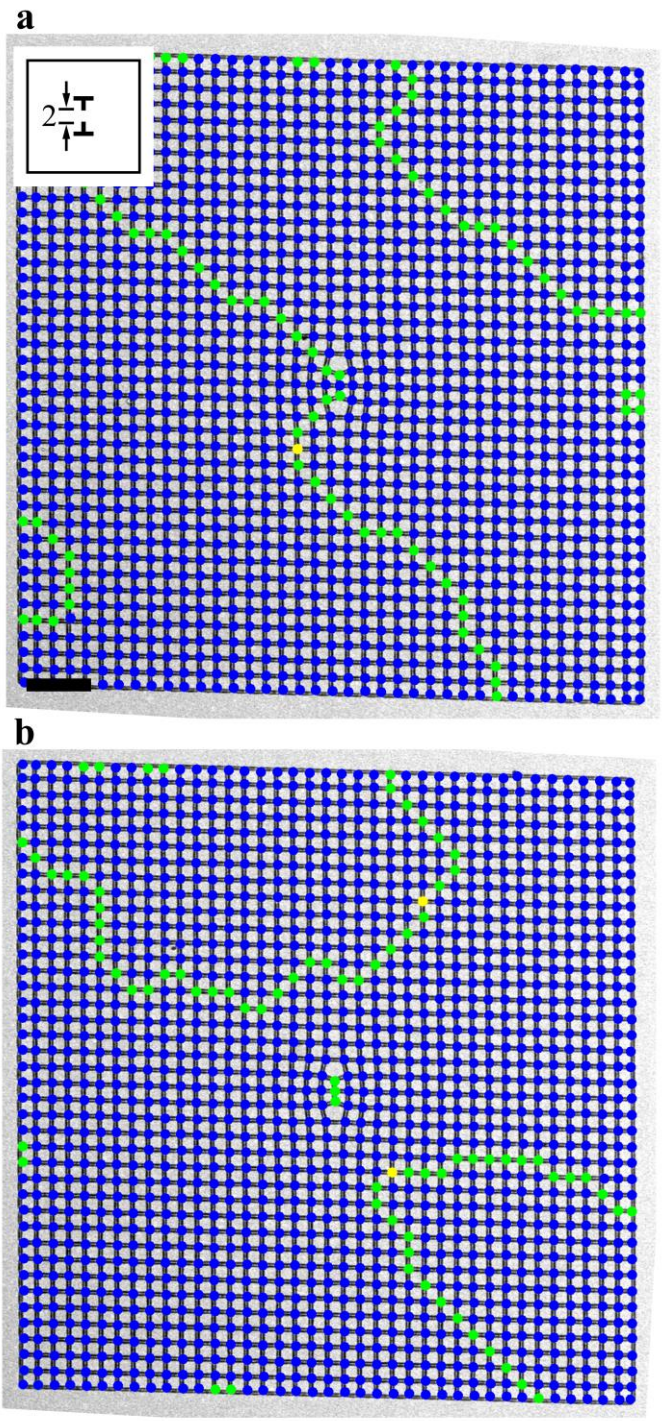

Figure 4.29 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *T*2 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

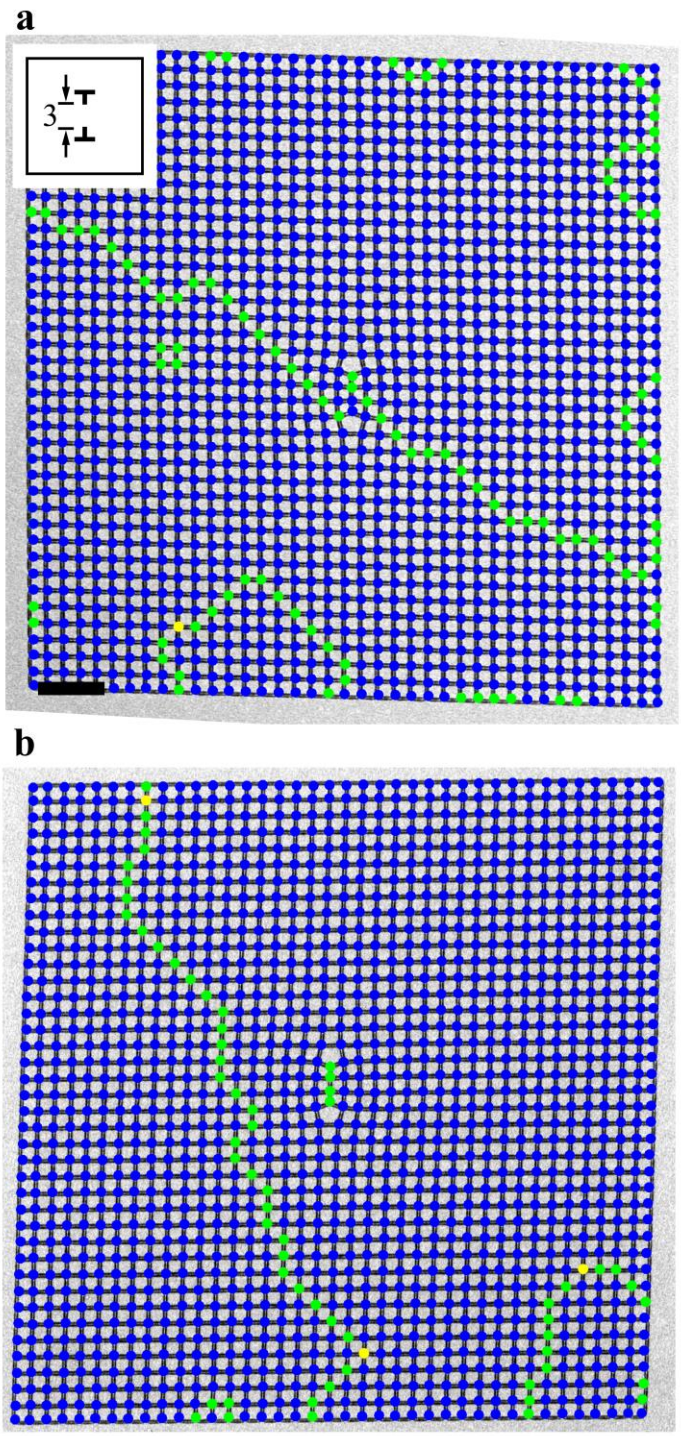

Figure 4.30 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *T*3 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

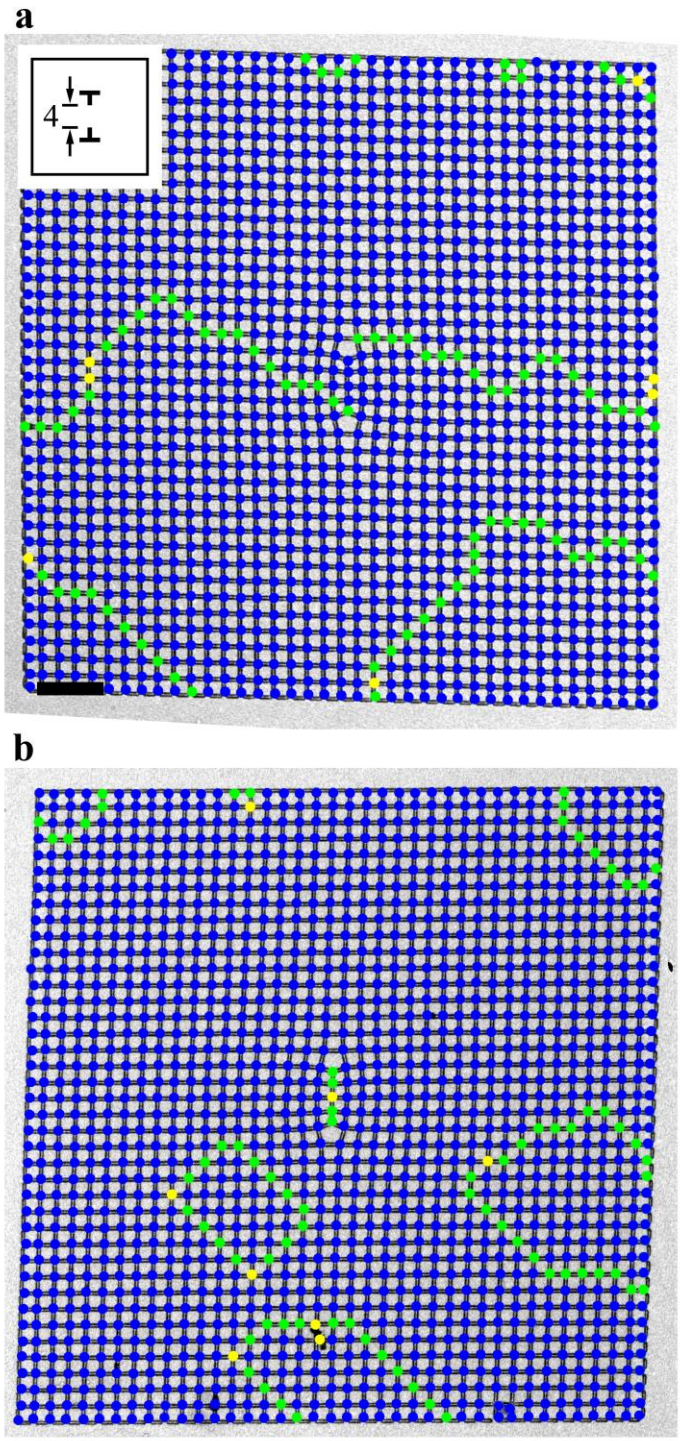

Figure 4.31 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *T*4 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

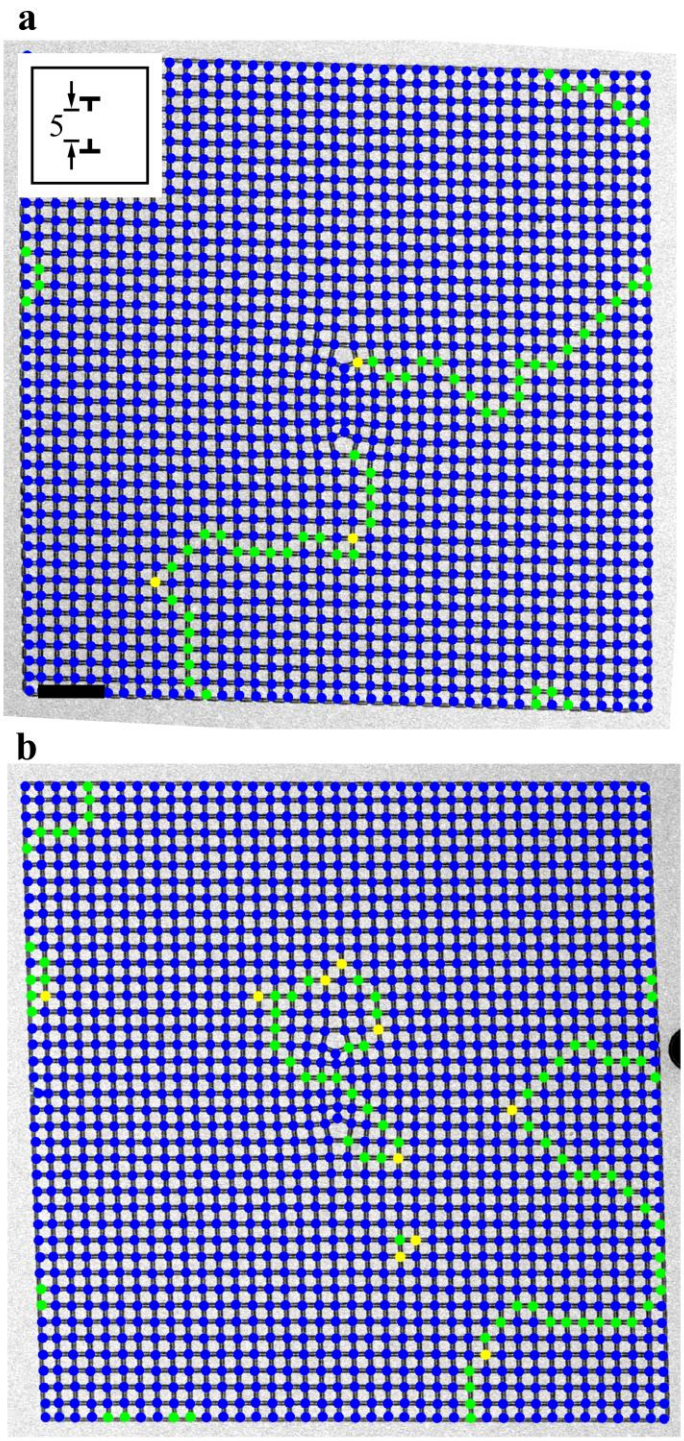

Figure 4.32 **Vertex maps of crystals with two topological defects.** Lorentz contrast TEM images with vertex types overlaid on *T*5 geometry crystals showing domain walls (**a**) not connecting and (**b**) connecting the dislocations. In (**a**), domain walls begin at each dislocation and end at the edges of the crystal. In (**b**), a single domain wall connects the two dislocations. Both configurations are very well ordered otherwise. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

From our observations, we compute the probability of domain walls not connecting dislocations in our samples for 30 different two-dislocation geometries and plot the probability vs. the distance between the two dislocations in Figure 4.33. Figure 4.33(a) shows this probability and Figure 4.33(b) has the same probability information, but also includes labels for each point specifying the geometries that contributed to that point. Note that some points have multiple geometries averaged together. Overall, we find that the probability for a domain wall not connecting the two defects is typically almost 1 for dislocations that are spaced greater than 10 lattice constants apart. Between 5 and 10 lattice constants, the probability hovers around 0.8, which is still quite high. At 5 lattice constants, the probability starts to drop sharply and reaches about 0.3 at one lattice constant.

We also examine the relative behavior of the different families of *T*, *A*, and *R* geometries. Figure 4.34 shows a subset of the data given in Figure 4.33 divided into the three main families. For larger lattice spacings, the three geometries all seem to have very similar behaviors and probabilities. A close examination of Figure 4.34(b) shows at the very shortest inter-defect distances there are signatures of a difference between the *T* and *A* geometries. These are crystals like those in Figures 4.20 and 4.29 for the *T* geometries and Figures 4.18 and 4.22 for the *A* geometries. The *T* geometries tend to have a higher rate of a domain wall connecting the two dislocations. It is interesting that when a domain wall connects two dislocations in a *T* geometry, a lot of the time the wall is a very straight line of type II vertices that goes directly from one defect to the other via the three-island vertex of the dislocation point. This is seen in Figures 4.20, 4.29, and 4.30. For the *A* geometries, the domain walls are much more likely to extend outwards from the three-island vertex as seen in Figures 4.18(a), 4.22, and 4.23(b). Because the walls from one

dislocation tend to form and extend away from the other dislocation in the *A* geometries, it is slightly less likely for a wall to connect the two dislocations in this geometry. There may be similar effects in the *R* geometries for the same reason discussed here, but it is harder to make a direct comparison because the defects cannot be as close together in the *R* geometry. Examples of the *R* family behavior are given in Figures 4.25, 4.26, 4.27, and 4.28. Here, it is again easy for a very short type II vertex chain to connect the two dislocations when the chain extends out of a three-island vertex at the defect point.

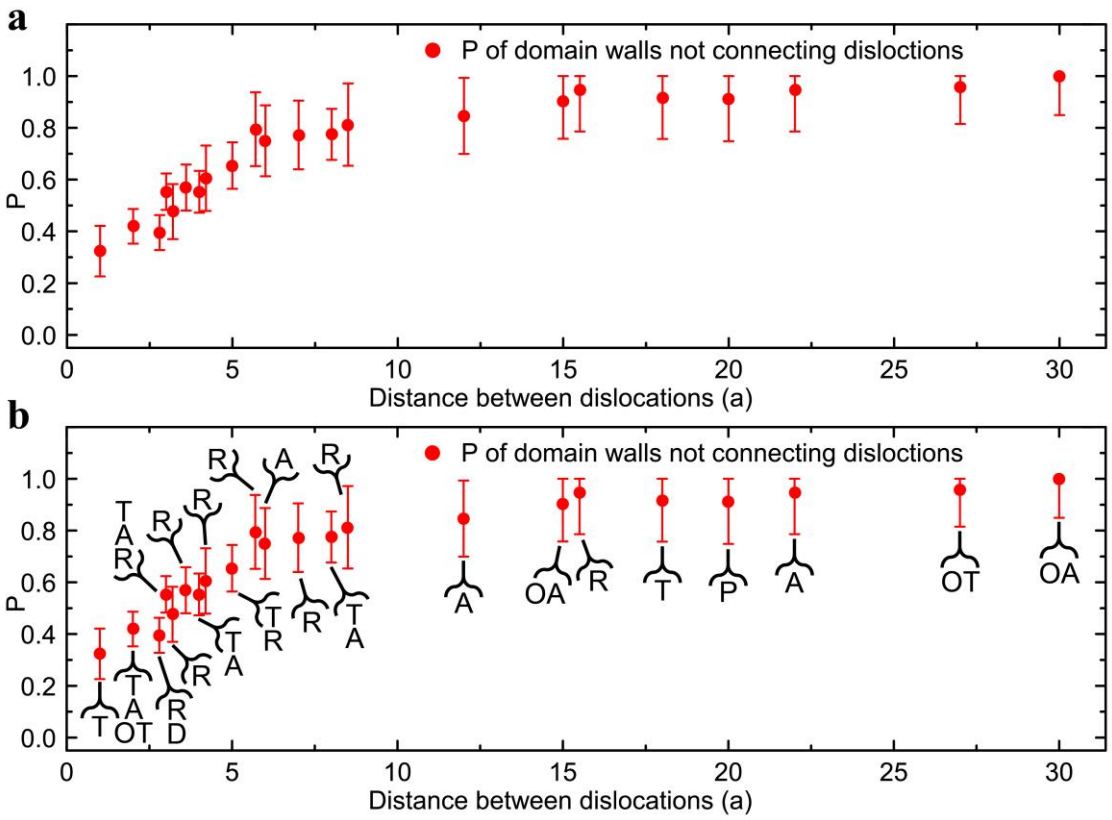

Figure 4.33 **Probability of domain walls not connecting dislocations vs. distance between dislocations in crystals with two topological defects.** Parts (**a**) and (**b**) present the same information, but in (**b**), letters are added indicating the geometries of the data points. Some data points include two or three geometries of equal distance apart. Distances are expressed in units of the lattice constant a ≈ 500 nm. Error bars are one standard deviation and are calculated from counting statistics. Each geometry includes approximately 40 crystals.

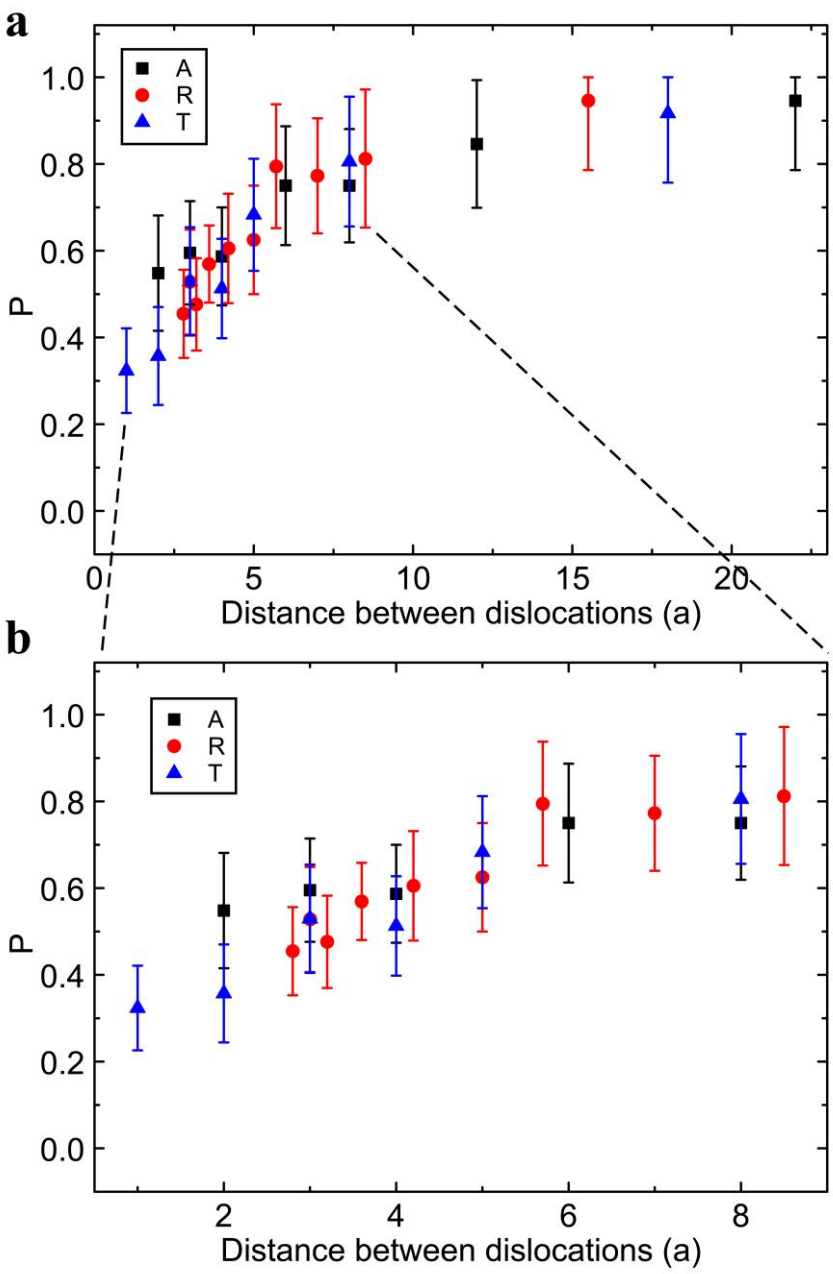

Figure 4.34 **Probability of domain walls not connecting dislocations for different geometry families.** Similar data to Figure 4.33, but probability vs. distance divided into the different families of geometries for direct comparison. Part (**b**) is a close up of a section of the full dataset in (**a**). Distances are expressed in units of the lattice constant a ≈ 500 nm. Error bars are one standard deviation and are calculated from counting statistics. Each geometry includes approximately 40 crystals.

Overall, we find the presence of a domain wall at a dislocation to be extremely robust. We studied five samples fabricated from different depositions of $FePd_3$ and collected over 1600 images of annealed crystals, and in every image, there is a domain wall originating from every dislocation and connecting to either the edge of the crystal or to another dislocation. For samples with two dislocations spaced more than 5 lattice constants apart, it is possible, but rare, for a single domain wall to connect both dislocations. It is much more common for the defects to behave independently and nucleate their own domain walls which extend throughout the system and rarely take the shortest available path to an edge of the crystal. We note that topological frustration, while it derives from crystalline disorder, is fundamentally different from standard disorder-induced frustration. Here, the system shows extended frustration as the topology of the defect causes long-range effects in the lattice. In 2D ASI, the dislocation is a 0D point defect, but it promotes itself into a 1D line of frustration. In a 3D material, dislocations are 1D defects, and this would result in an extended 2D surface of frustration of the system.

4.2.2 Theoretical framework and discussion

In this section we discuss our results in the context of crystalline materials systems and develop a theoretical framework for the presence of topological frustration within an ordered system. An important question remains about how our observations relate to dislocations in 3D crystalline materials and the outstanding issue of whether a unique ground state could be realized in spin ice. We consider previous measurements of residual entropy in spin ice in the context of sample preparation and dislocation density. The

material studied in the seminal report by Ramirez *et al*. [13] was a pressed powder sample, which should result in a high density of dislocations both due to damage during compaction and also due to grain boundaries in the resulting structure. This material, with a presumably high density of dislocations, showed a level of entropy equal to the Pauling ice value. On the other hand, the $2 \times 1 \times 1$ mm^3 single crystal studied by Pomaranski and coworkers showed less entropy, a small amount below the Pauling value [57]. Pomaranski *et al.* also studied powder samples, which showed the same qualitative behavior as the samples in Ref. [13]. We suspect that the topological frustration due to a high density of dislocations in both powder samples could be preventing ground-state ordering, while the single crystal in Ref. [57] could be showing a slight emergence of the ground state because it contains fewer topological defects. However, it is important to note that larger single crystals maintain lattice coherence over greater distances and are thus actually more susceptible to topological frustration. As an example, a cubic crystal of side length N unit cells and volume N^3 unit cells could need only on the order of $N_d = N$ dislocations to be fully frustrated, if the emanating frustrated surfaces are non-interacting, similar to the domain walls we observe in our specimens. The dislocation density would be $\rho_d = N_d/N^2$, so the density required to frustrate a system scales as $1/N$, implying that larger single crystals are actually more likely to be impacted by the presence of dislocations than smaller crystals. This is counter to the conventional expectation that larger single crystals show more intrinsic ground-state behavior. On the contrary, smaller single crystals, or single crystals with lower densities of dislocations, should be better able to reach a fully ordered ground state when topological frustration is present.

Implications of topological frustration could extend far beyond the systems discussed above, and we now lay out a framework for such considerations as follows. Many materials, simple and exotic alike, display ordered states that can generally be characterized by an order parameter and associated ordering vector $\mathbf{q} \neq \mathbf{0}$. In our square ASI, there are two q vectors corresponding to the doubly degenerate ground state, $\mathbf{q} = \pm(\pi/a)[11]$, which tile type I vertices onto the lattice. Using these q vectors, spin directions in the ordered states can be assigned through a Bloch form

$$\boldsymbol{S}_i = u_i(\boldsymbol{r})e^{i\boldsymbol{q}\cdot\boldsymbol{r}} \tag{4.2}$$

where **q** is the ordering vector, **r** is the coordinate of the spin, i refers to a given sublattice, either 1 or 2, and $u_i(\mathbf{r})$ acts like a key, specific to the choice of basis and origin. The q vectors depend on the choice of basis and origin as well. For $\mathbf{q} = \pm(\pi/a)[11]$, we have chosen the origin at the middle of a vertex where 4 elements meet. The system has a square Bravais lattice with lattice constant *a* and we choose the basis vectors for the two spins per unit cell as [a/2 0] and [0 –a/2]. For the horizontal sublattice, which we will call sublattice 1, $u_1(\mathbf{r})$ specifies that a spin labeled as an "*i*" is in the +x direction and a spin labeled as a "-*i*" is in the –x direction. For the vertical sublattice, sublattice 2, $u_2(\mathbf{r})$ specifies that a spin labeled as a "-*i*" is in the +y direction and a spin labeled as an "*i*" is in the –y direction. These definitions completely specify the two possible ground states of square ASI with $\mathbf{q} = \pm(\pi/a)[11]$ as the ordering vectors, with each vector corresponding to one of the two degenerate ground states. As an example using the positive q vector, consider two neighboring spins on sublattice 2 located at (0, a/2) and (a, a/2). If we compute **S** for each spin, we find that the first one at (0 , a/2) is assigned an *i* and the other spin is assigned a -*i* indicating that they are pointing in opposite directions, consistent with ground-state

ordering in the system. If we choose the other q vector, the values assigned to the two spins both switch, so they remain pointing in opposite directions. It is straightforward to show that this continues for both sublattices throughout the whole lattice.

We can now state our theoretical description of topological frustration: If the topological Burgers vector of a defect is perpendicular to a given **q**, or contains a complete wave of the order parameter with $\mathbf{q}\cdot\mathbf{b} = 2\pi n$, where n = 0 or an integer, the system will not be frustrated. However, if these conditions do not hold, the system will exhibit topological frustration. To demonstrate this idea we now consider ordering in the square ASI described above. Figure 4.35 shows a schematic of a coordinate system with a dislocation. There is a spin of interest at location **r**. We can take a path **x** + **y** (as defined in Figure 4.35) to the spin, or alternately, we can take an equivalent path **k** + **l** + **m** + **n** to the same spin. However, if the second path encloses a topological defect, we will need an extra vector, the Burgers vector **b**, to reach the same spin. Now, when we compute **q·r** for the two paths, we should get the same result because we are specifying the same spin in the ordered lattice. Since **x + y = k + l + m + n**, these terms dotted with **q** give the same answer each time, but there is an additional **q·b** term for the second path. If $\mathbf{q}\cdot\mathbf{b} = 2\pi n$, $e^{i(\mathbf{q}\cdot\mathbf{b})} = e^{i(2\pi n)} = 1$, and the two paths have the same result for the given spin and there is no issue, but if **q·b** equals anything else, there is a difference and the system is frustrated. If the second path encloses two topological defects with Burgers vectors that add to 0, there will also not be difference between the two paths because any additional terms will cancel out.

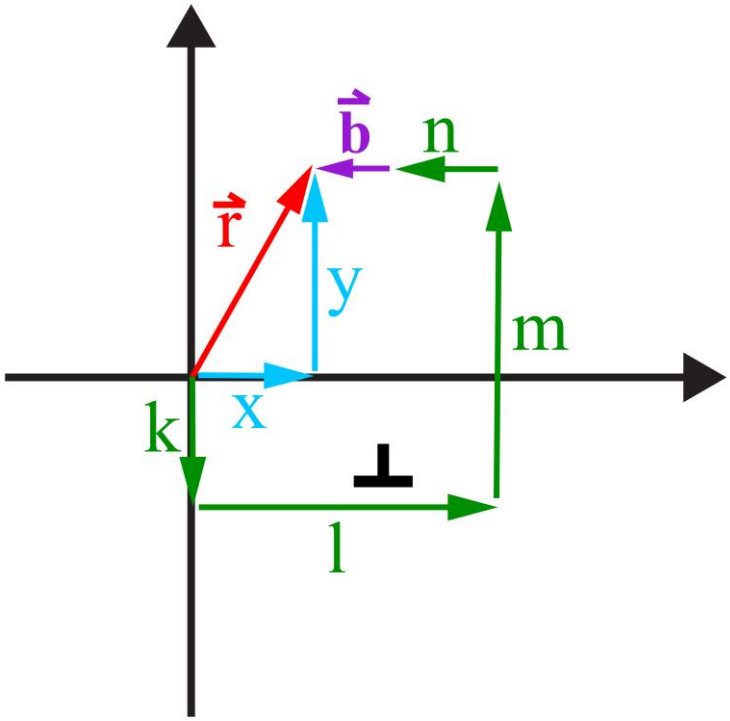

Figure 4.35 **Multiple paths near a topological defect**. Schematic showing how multiple paths can be taken to reach the same point in 2D space. Here, $\vec{r}$ represents the given point. $\vec{x} + \vec{y} = \vec{k} + \vec{l} + \vec{m} + \vec{n}$, but the green path requires an extra vector, the Burgers vector $\vec{b}$, to reach the point specified by r when it encloses a dislocation.

All the geometries we have studied thus far have had a Burgers vector in the [10] or [01] direction and these will always be non-orthogonal to ordering vectors $\mathbf{q} = \pm(\pi/a)[11]$ in square ASI, thus the system will always be topologically frustrated. Now, we demonstrate that a sample with a Burgers vector in the [11] direction will not be topologically frustrated. We have fabricated crystals with dislocations that have Burgers vectors in the [11] direction using the same methods described in this chapter. An in-focus image of a **b** = a[11] topological defect is shown in Figure 4.36 with the Burgers loop overlaid confirming the orientation of the Burgers vector. I will refer to crystals of this geometry as *B*11. A full *B*11 crystal is shown in Figure 4.37.

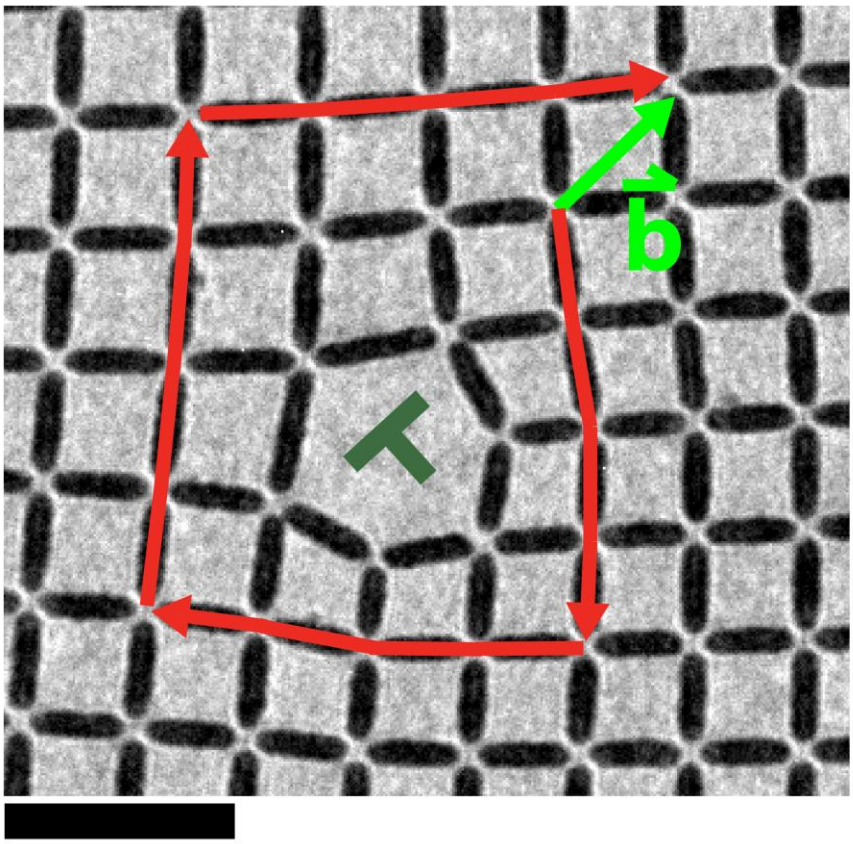

Figure 4.36 **[11] Burgers vector topological defect**. In-focus TEM image of a section of a crystal containing an edge dislocation with Burgers vector **b** = a[11]. The Burgers circuit and vector are overlaid in the image confirming the topological nature of the defect. Scale bar is 1 μm.

We anneal and image the crystals in the same way as the other experiments in this chapter. When we do, we find that the *B*11 crystals indeed do not require a domain wall originating from the dislocation point and the crystal can support continuous ground-state ordering around a defect of this nature. A vertex map around a **b** = a[11] dislocation is given in Figure 4.38 and full crystals containing one of these defects are given in Figures 4.39 and 4.40. It is interesting that in Figure 4.40, there are two domain walls that come out of the dislocation point. The crystal here is not topologically frustrated, so it cannot have exactly one domain wall originating from it, but it can have two. These images give experimental verification of the theoretical framework for our system.

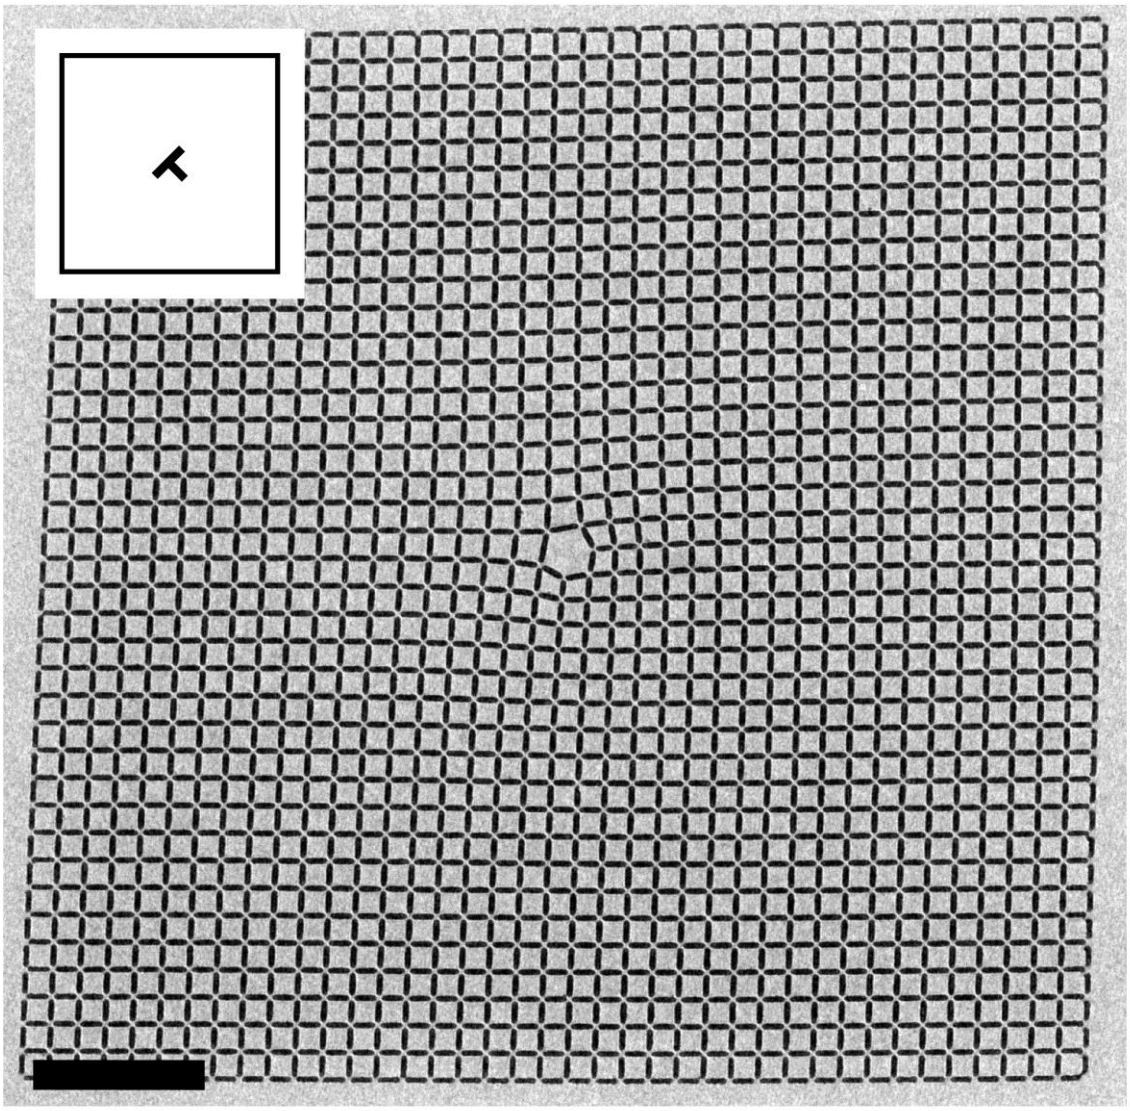

Figure 4.37 **Full crystal with a [11] Burgers vector topological defect**. In-focus TEM image of a full crystal showing one topological defect with a Burgers vector **b** = a[11], referred to here as a *B*11 geometry crystal. Inset shows the schematic representation of this crystal. Scale bar is 3 μm.

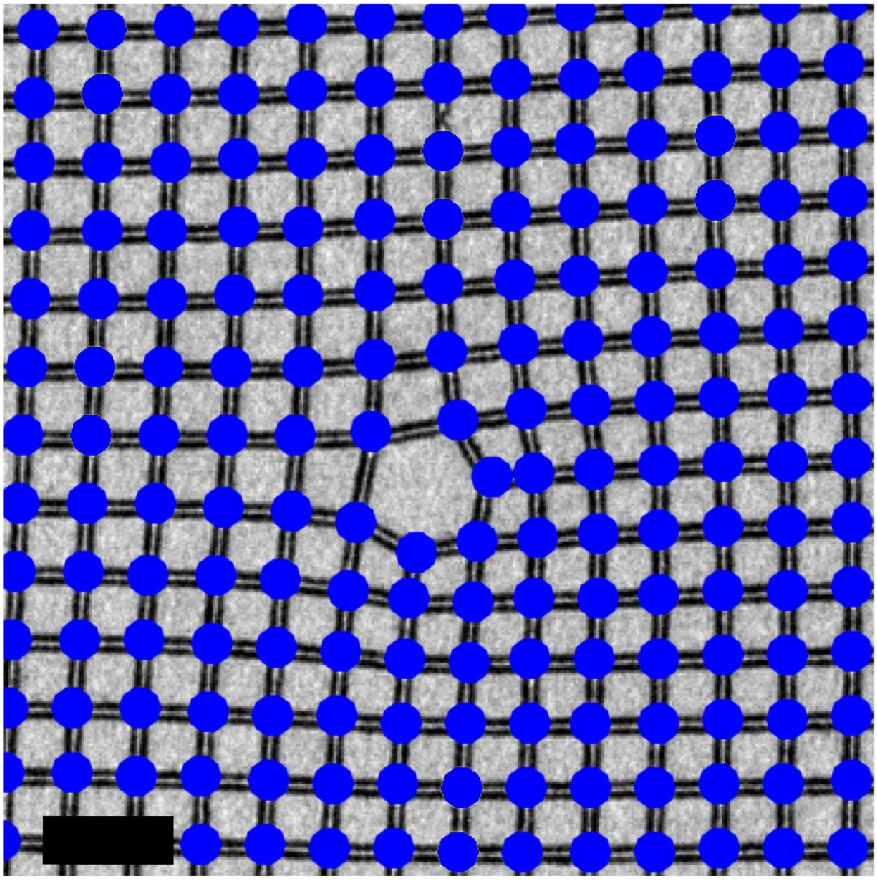

Figure 4.38 **Vertex map of a [11] Burgers vector topological defect.** Lorentz TEM image with vertex types overlaid on a section of a crystal containing an edge dislocation with Burgers vector **b** = [11]. This is a topological defect, but because of the Burgers vector orientation, it can still support perfect ground-state ordering around it. Scale bar is 1 μm.

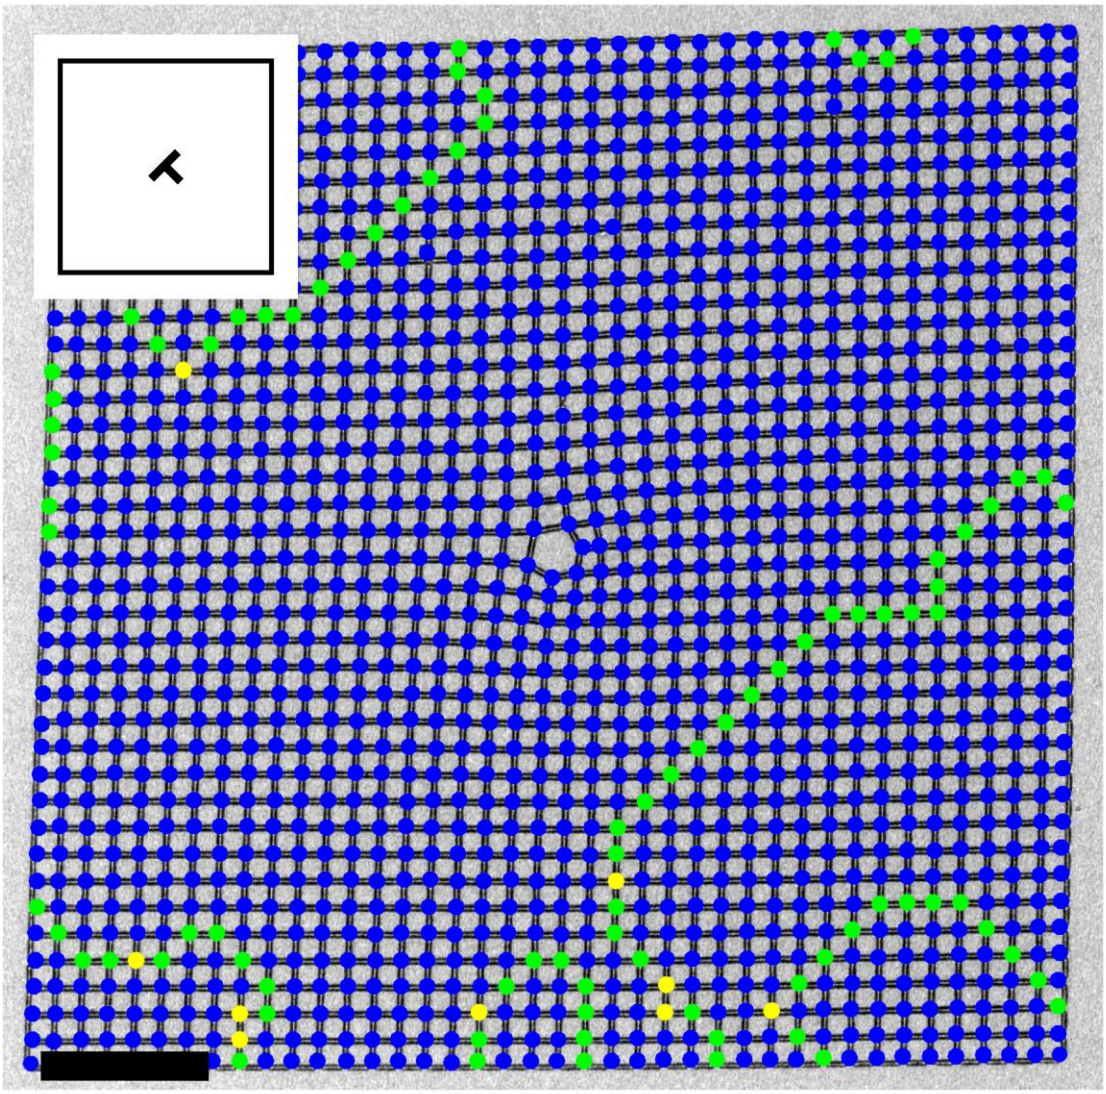

Figure 4.39 **Vertex map of a [11] Burgers vector topological defect full crystal.** Lorentz TEM image with vertex types overlaid of a full crystal containing an edge dislocation with Burgers vector **b** = a[11]. The defect does not nucleate a domain wall, though other traditional square ASI domains and domain walls are present. Inset shows the schematic representation of this crystal. Scale bar is 3 μm.

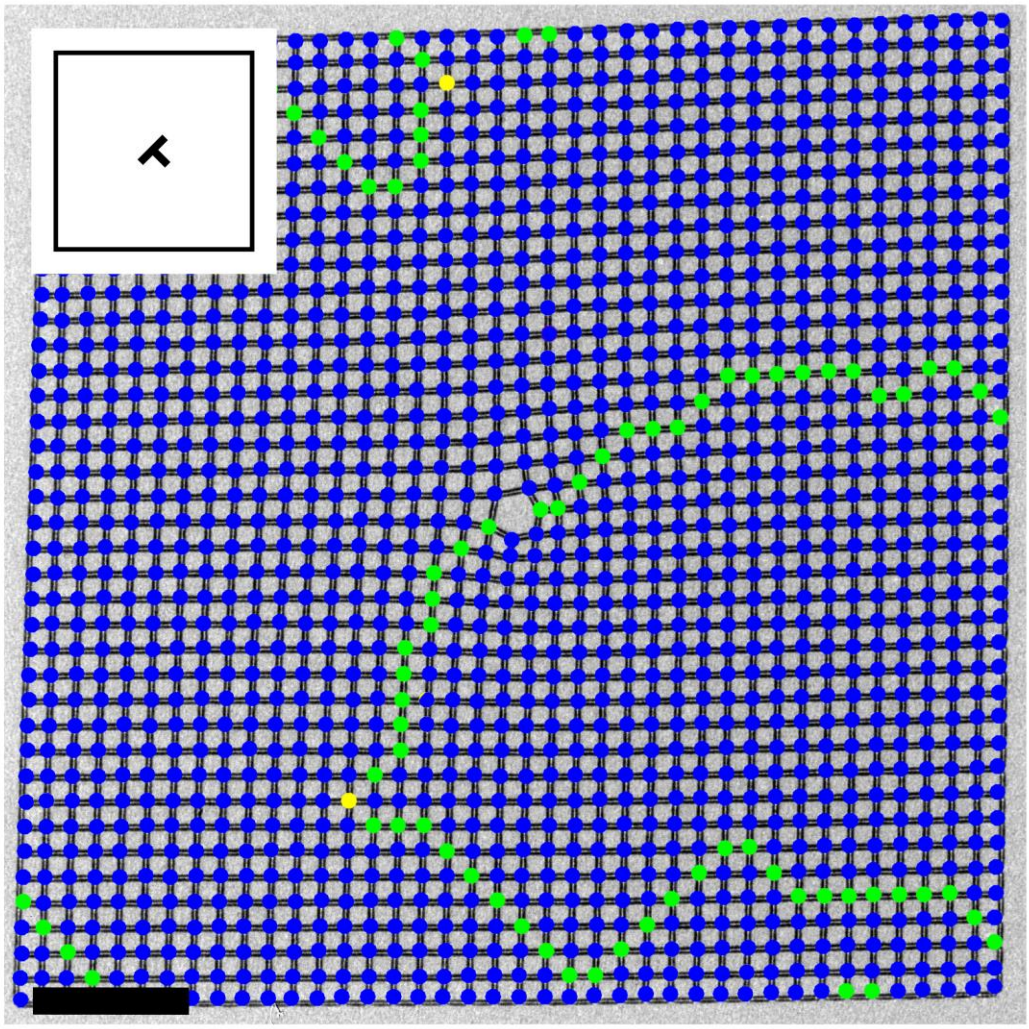

Figure 4.40 **Vertex map of a [11] Burgers vector topological defect full crystal.** Lorentz TEM image with vertex types overlaid of a full crystal containing an edge dislocation with Burgers vector **b** = a[11]. The defect cannot nucleate a single domain wall, but can support two domain walls intersecting it. Inset shows the schematic representation of this crystal. Scale bar is 3 μm.

Returning now to the broad and general impact of topological frustration on 3D crystalline systems with nonzero **q** vectors, we consider some concrete examples. In the pyrochlore spin ices, there are 6 **q** vectors for 6 realizations of the predicted long-range ordered state, **q** = (2π/a) <100> [51]. A schematic of the ground state with **q** = (2π/a)[001] is shown in Figure 4.41. If we assume Burgers vectors of dislocations to point along one of the primitive FCC lattice vectors, (a/2) <110> , then, for each possible Burgers direction,

4 out of the 6 ground-state **q** vectors will not be perpendicular to the Burgers vector and will thus be frustrated. This is shown schematically in Figure 4.42 with a Burgers vector in the $(a/2)[10\bar{1}]$ direction as an example. This Burgers vector will not cause any structural mismatch away from the dislocation core [214], but it will disrupt the ground-state ordering in the lattice. The dislocation creates a domain wall in the $\mathbf{q} = (2\pi/a)[001]$ ordering of the system. The location of this domain wall is arbitrary, but it must be present. It is easy to calculate that $\mathbf{q}\cdot\mathbf{b} \neq 2\pi n$, thus the system is topologically frustrated.

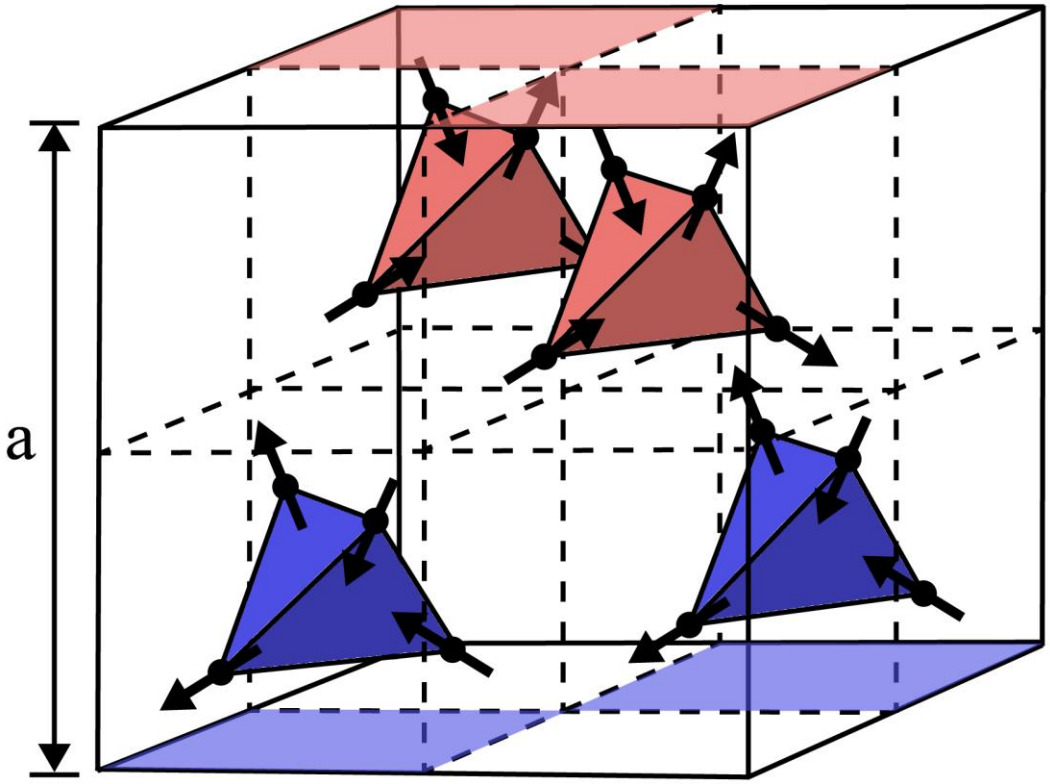

Figure 4.41 **Pyrochlore spin ice ground state with order parameter q = (2π/a)[001].** Reduced unit cell showing the locations of the rare earth atoms and their magnetic moment directions in the **q** = (2π/a)[001] ground state. This state has alternating planes of ordered tetrahedra, here colored red and blue, corresponding to the planes of ordering in Figure 4.42.

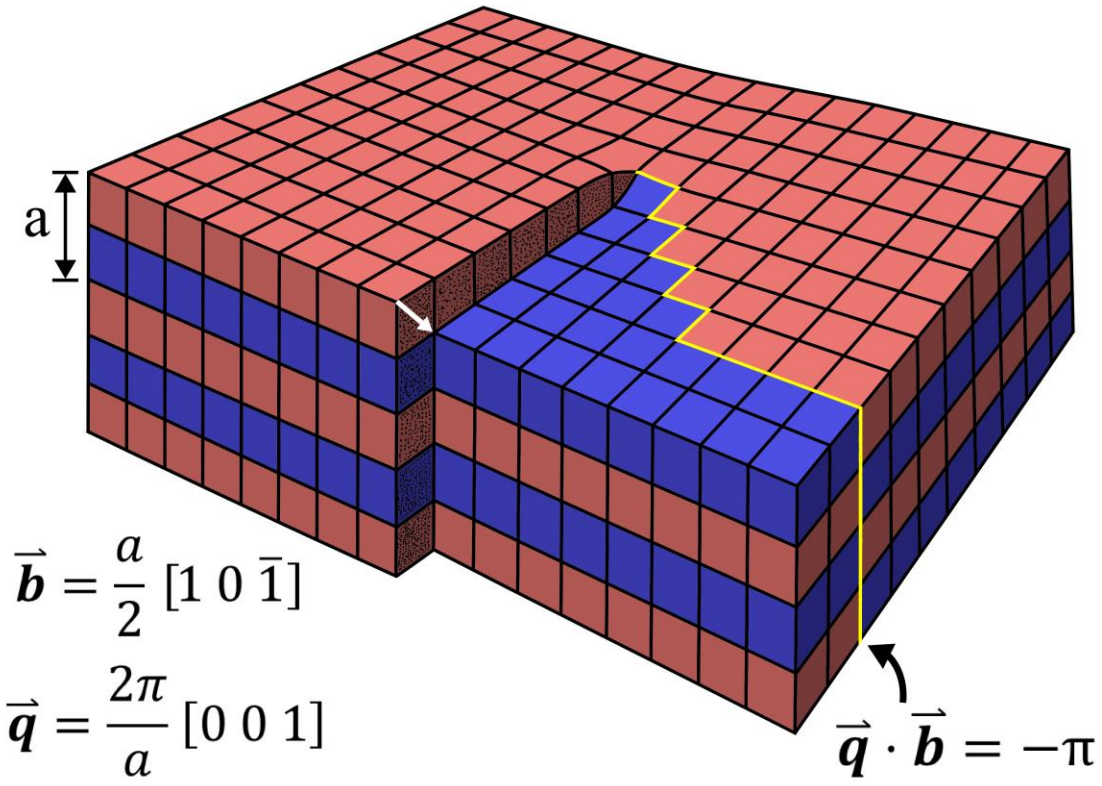

Figure 4.42 **Topological frustration in pyrochlore spin ice**. Schematic showing how a dislocation in pyrochlore spin ice could disrupt ground-state ordering. Each cube is 1/8 the pyrochlore unit cell shown in Figure 4.41. This example has a dislocation with a Burgers vector (shown as a white arrow) in the (a/2)[10$\bar{1}$] direction which is a primitive basis vector for the FCC parent structure of the pyrochlore lattice so that this does not cause any structural mismatch away from the dislocation. This dislocation does create a domain wall in the **q** = (2π/a)[001] ground-state order parameter, highlighted in yellow. The location of this domain wall is arbitrary, but it must be present. Note that **q**·**b** ≠ 2πn and thus the system is topologically frustrated.

If a pyrochlore crystal contains dislocations with all three FCC primitives, then all 6 ground states will exhibit frustration, and the specimen could be prevented from finding its unique ground state, as is observed experimentally. This frustration should be present in many other systems with **q** ≠ **0** order parameters. Prime examples may include spin-density-wave materials [215], such as antiferromagnets and ferrimagnets, antiferroelectrics, such as $SrTiO_3$ [216], charge-density-wave materials [217], and pair-

density-wave superconductors [218]. In cases where the dislocation densities are low, or where the frustrated domain walls have stronger surface tension between dislocations, the frustration may not prevent the formation of the order parameter, but rather it will only degrade it or delay it somewhat. In such cases, we expect topological frustration may thermally broaden the phase transition, degrading the formation of the $\mathbf{q} \neq \mathbf{0}$ order parameter. Topological defects may even play a role in the frustration of ice-XI, the thermodynamic ground state of water not known to form under typical experimental conditions [10]. In general, the effects of topological frustration could be observed in a wide variety of materials systems, and artificial spin ice could serve as a valuable platform for future studies to develop these general considerations.

We aim to argue here that dislocations should receive more attention in theoretical work and interpretations of experimental systems with $\mathbf{q} \neq \mathbf{0}$ than they currently do. A recent and extremely thorough study on $Dy_2Ti_2O_7$ spin ice by Henelius and coworkers indeed does not discuss the possibility of dislocations affecting the ordered states [219]. The paper does suggest, however, that "random disorder" is the key underlying cause of the varied and unique behavior observed in this compound. The authors speculate that either stuffed spins – extra magnetic atoms in the lattice – or oxygen vacancies should be responsible for the upturn in heat capacity and corresponding reduction of entropy observed by Pomaranski *et al.* [57,219]. There have been a number of studies on the effects of stuffed spins in pyrochlore spin ices [220–222] as well as oxygen vacancies [223]. Henelius *et al*. rule out most plausible reasons for the reduction of entropy. They try including higher order terms in the dipolar spin ice model and exploring more sophisticated models using advanced computational techniques. Through all this work, their main

conclusion is that the system is extremely sensitive to disorder. Based on our results with topologically frustrated ASI, we argue that dislocations in the crystals are just as likely to be a source of random disorder and cause the range of observed behaviors, based on the specific density of dislocations in the materials studied. Indeed, dislocations may be better able to account for the observed frustration, since 0D point defects only produce 0D regions of frustration, whereas 1D dislocations produce 2D sheets (domain walls) of frustrated regions (see Figure 4.42). Overall, there are still many unresolved questions about a ground-state transition in spin ice and more work is needed on higher purity samples and at lower temperatures. ASI presents a highly tunable system that will aid in these studies by providing unique information at the single spin level.

4.2.3 Additional observations

In this section we survey a wide range of experimental observations about our topological defect samples. We reference some figures already presented in this chapter as well as introduce some more rare and specific behaviors in a variety of different geometries. We begin by examining edge effects in our crystals. Figure 4.14 shows a common feature found in many different topological defect crystals and geometries. In the bottom right corner there is an isolated pair of type III vertices. If a single spin flips within a domain of ground-state ordered type I vertices, it will create a pair of type III vertices. These isolated defects do not greatly affect the overall behavior of the lattice, but do indicate some interesting ideas. They appear somewhat rarely, though some examples are also seen in Figures 4.19, 4.21, 4.46, and 4.48. At the top of Figure 4.14 is a pair of type II

vertices on the edge of the crystal. This type of defect is again caused by flipping a single spin among ground-state order, but these defects located at the edges are much more common. This is because the edge vertices have only three elements and the energy difference between the type I and type II three-element vertices is much smaller than the corresponding energy difference for 4-element vertices. Edge elements also tended to come out slightly thinner than bulk elements because of proximity effects during lithography. This further reduces the intrinsic energy barrier for these elements to spontaneously flip into energetically unfavorable arrangements. Type II vertex pair defects on the edges of crystals are very common and are seen in many of the figures in this chapter. Some examples include Figures 4.20, 4.24, and 4.27. Figure 4.39 has a chain of 4 type II defects in a row on the left edge of the crystal. There are many examples of small domains that are contained very close to an edge, likely because it is easy to nucleate them close to the edge, but more energetically unfavorable for the domains to grow further into the lattice. Some of these small domains can be seen, for example, in Figures 4.22, 4.23, 4.24, and 4.30.

We have discussed topological defects at length in this chapter, but it is important to think about other types of defects as well. We note that a vacancy, a single missing magnet without lattice distortion, is not a topological defect and does not induce frustration or nucleate a domain wall; a square ASI lattice can still display perfect order with a single spin removed. This is seen in Figure 4.43 for a simple vacancy. Here, the missing spin is an unintentional defect from fabrication, but it illustrates an important point. We also intentionally fabricated crystals with a vacancy, but allowed the system to relax in the same way we created the topological defect lattices. This crystal is shown in Figure 4.44 and can be thought of an *A*1 geometry, though it was not included in Figures 4.33 or 4.34 because

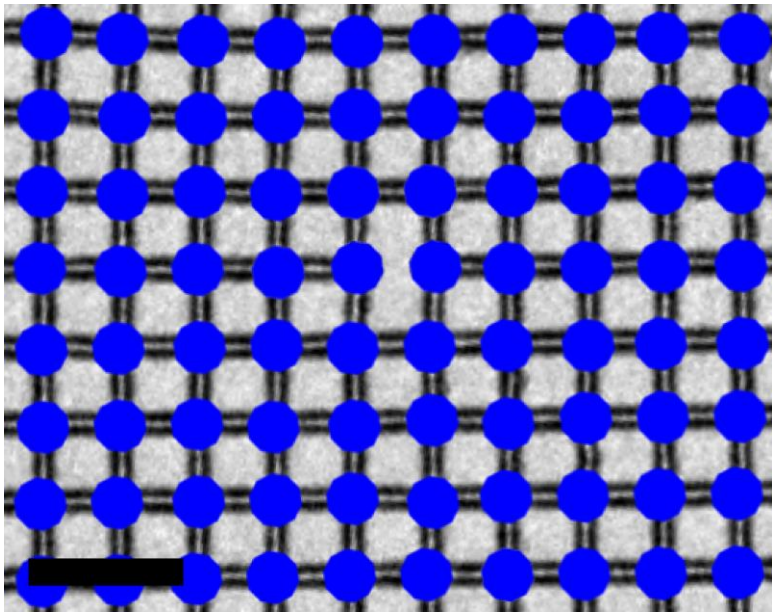

Figure 4.43 **Vertex map of a non-topological defect.** Lorentz TEM image with vertex types overlaid on a square geometry sample with a vacancy. A simple vacancy, a single missing spin in an otherwise perfect square lattice, does not nucleate a vertex chain or result in topological frustration. Long-range ground-state ordering is not affected by the presence of a vacancy since it is not a topological defect. Scale bar is 1 μm.

it is not a true topological defect. Still, this crystal displays two main behaviors shown in Figure 4.44. The lattice can either maintain continuous ground-state order around the non-topological defect as in Figure 4.44(a), or the defect can have two domain walls coming out of it, as in Figure 4.44(b). This situation is very similar to the *B*11 geometry crystals shown in Figures 4.39 and 4.40. These examples are interesting because it shows that the defects can nucleate domain walls even when they are not required and they do seem to nucleate them more frequently than would occur in a randomly selected defect-free location in one of these crystals. We also find it very interesting that when we compute, from our data, the probability of having two domain walls originating from the non-topological defects vs zero domain walls, for both the *A*1 and *B*11 geometries we get a probability of about 0.5. This is quite similar to the probabilities for the closest spaced

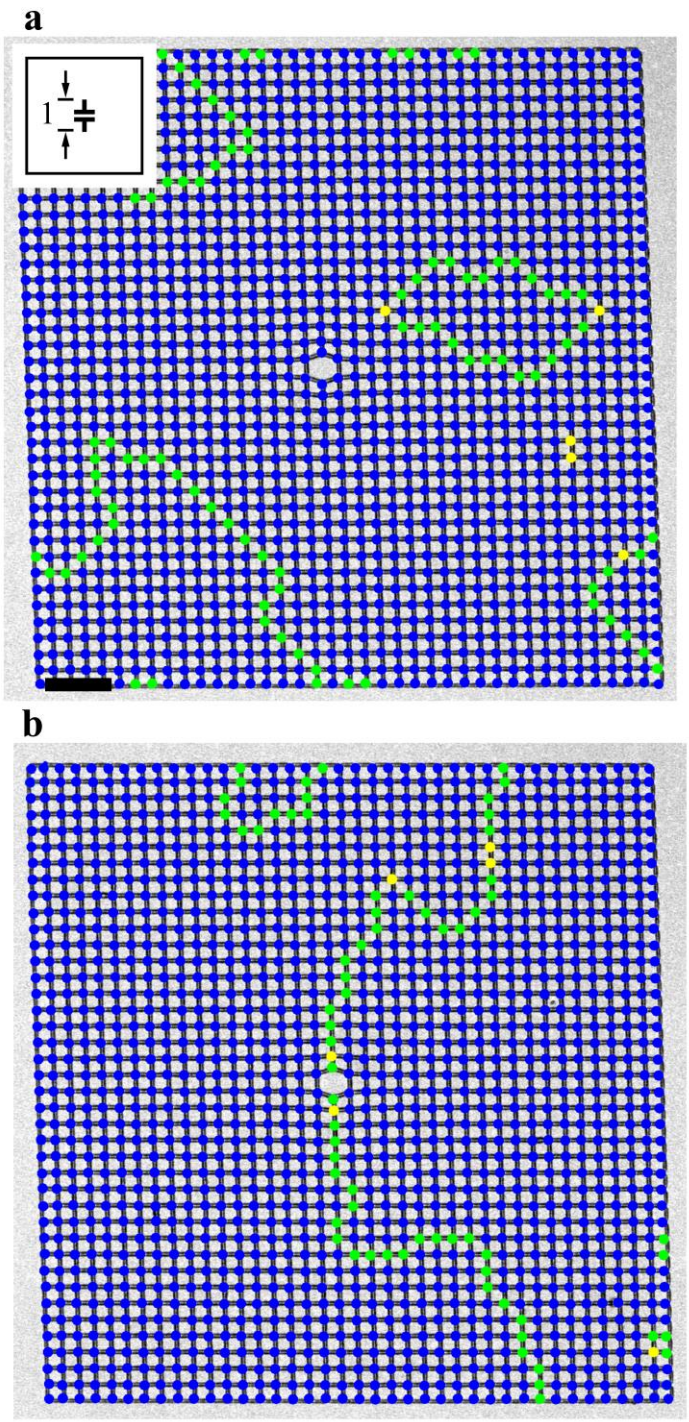

Figure 4.44 **Vertex map of non-topological defect crystals.** Lorentz contrast TEM images with vertex types overlaid on *A*1 geometry crystals with a vacancy. In (**a)**, there is perfect ground-state ordering surrounding the defect. In (**b)**, two domain walls originate from the defect point and extend to the edge of the crystal. Both configurations are very well ordered otherwise. Inset shows the location of the effective dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

two-dislocation crystals to have two independent domain walls originating from the individual defects.

We know that disorder plays an important role in our samples from the kagome crystals in Chapter 3 and the canonical square ASI samples in Chapter 2. Indeed, we will discuss this more in the next section, but I will highlight a few observations about disorder here. Figure 4.45 shows the same crystal that has been annealed twice and imaged after each time. We see that there is a completely different domain pattern in the two crystals. The behavior is qualitatively similar, but the same quenched disorder in the system gives different low temperature configurations after annealing. We have imaged almost every crystal we made multiple times (typically 2-3 times) and we have never seen a completely reproduced domain pattern in any sample. Figure 4.46 shows an in-focus TEM image along with a corresponding Lorentz contrast image of the same crystal. In this figure, we see how point defects in the individual elements, here as elements that came out too wide, can affect the domain wall patterns. We see that some of the defects, highlighted by the red circles, can act as pinning sites for a domain wall. We note, however, that these point defects are not required to pin a wall — as the topological lattice defects are — as the two defects highlighted in green do not have a wall associated with them. This shows that these fabrication defects can contribute to domain wall pinning, but they don't always have this effect.

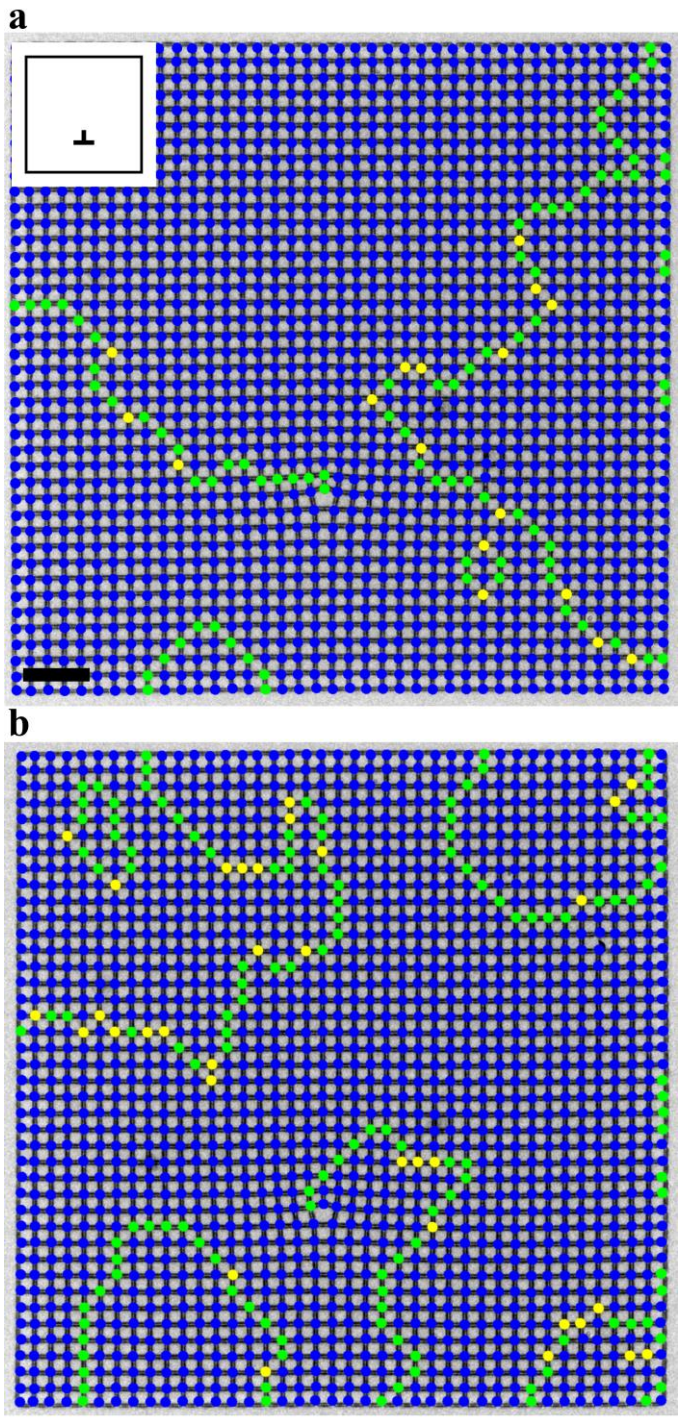

Figure 4.45 **Repeated annealing of the same crystal.** Lorentz TEM images with vertex types overlaid on the same 1*E* geometry crystal after repeated annealing. The same crystal annealed twice under the same conditions results in different domain configurations. The quenched disorder in the sample gives qualitatively similar, but not exactly reproducible ordering of vertices. Scale bar is 2 μm.

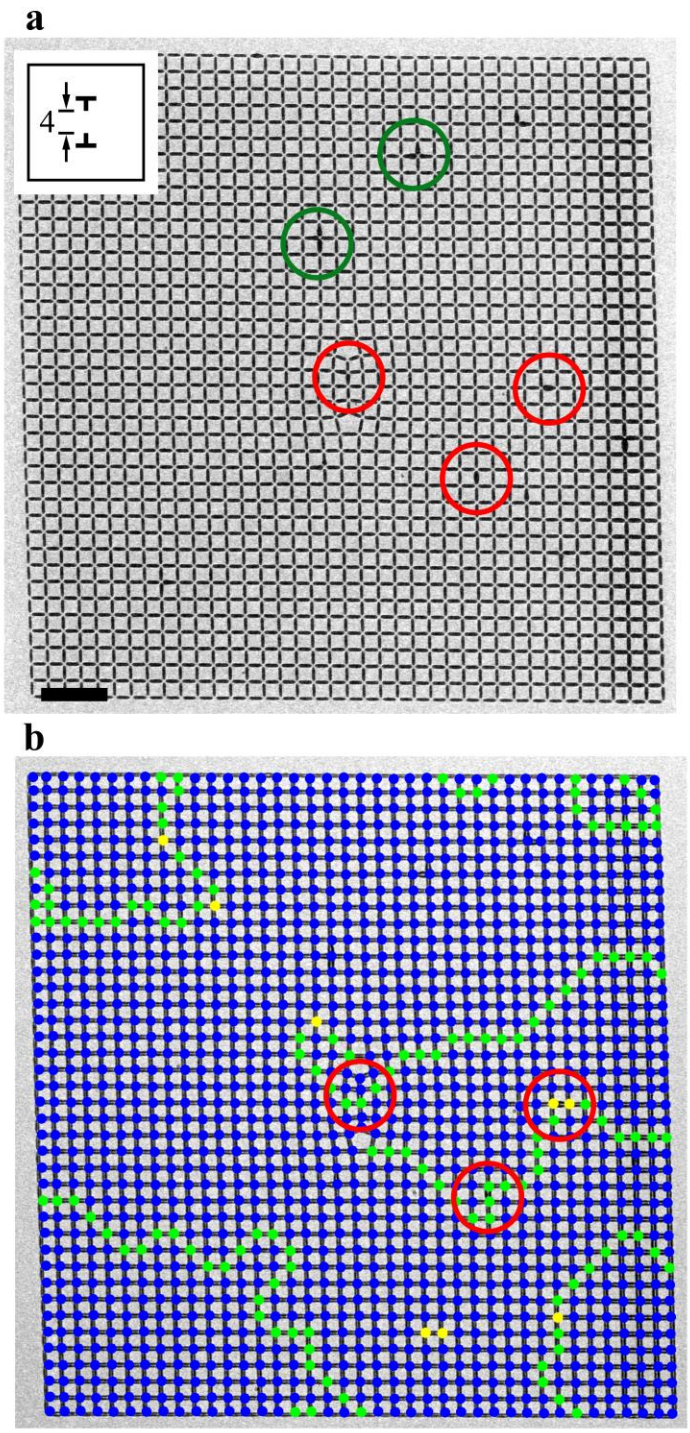

Figure 4.46 **Domain wall pinning sites**. **a**. In-focus TEM image of a *T*4 geometry crystal. **b**. Lorentz Contrast TEM image with vertex types overlaid of the same crystal as in (**a**). Multiple point defects are highlighted with red and green circles in (**a**). Three of these defects, highlighted in red, appear to pin a domain wall in (**b**), while the other two defects, highlighted in green, do not. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

We see a range of behaviors in the domain walls of our crystals. We have already discussed this idea briefly, but it should be restated that the walls rarely take the shortest path to an edge of the crystal. Figure 4.47 shows an example of some long, meandering domain walls. While this behavior is slightly more common, we do see some very straight and short domain walls as well, shown in Figure 4.48. Both of these crystals are *R* geometries, but have completely different domain wall lengths among otherwise nearly perfectly ordered crystals. These behaviors are common throughout all of our experimental data.

Topological defects require one domain wall originate from them, but similar to the observation that non-topological defects can have zero or two domain walls associated with them, the topological defects can also support three domain walls. This behavior is rare in our experimental samples, but Figures 4.49 and 4.50 show some unique examples. In Figure 4.49, one defect has a domain wall that connects to the other defect, but then that second defect has two more walls originating from it as well. In Figure 4.50, one of the defects has a single wall that runs to an edge of the crystal and the other defect has three walls that originate from it and all individually run to the edges of the crystal. Both of these configurations are extremely rare in experiment. For reasons not yet fully understood, this behavior was actually much more prevalent in our simulations, as discussed in the next section. In the experimental samples, we only saw three domain walls from one defect a handful of times and generally this occurred in lower quality crystals with more disorder. This suggests that the three domain wall configuration is energetically unfavorable and requires a sample with higher disorder, far from the ground state, to support the phenomenon.

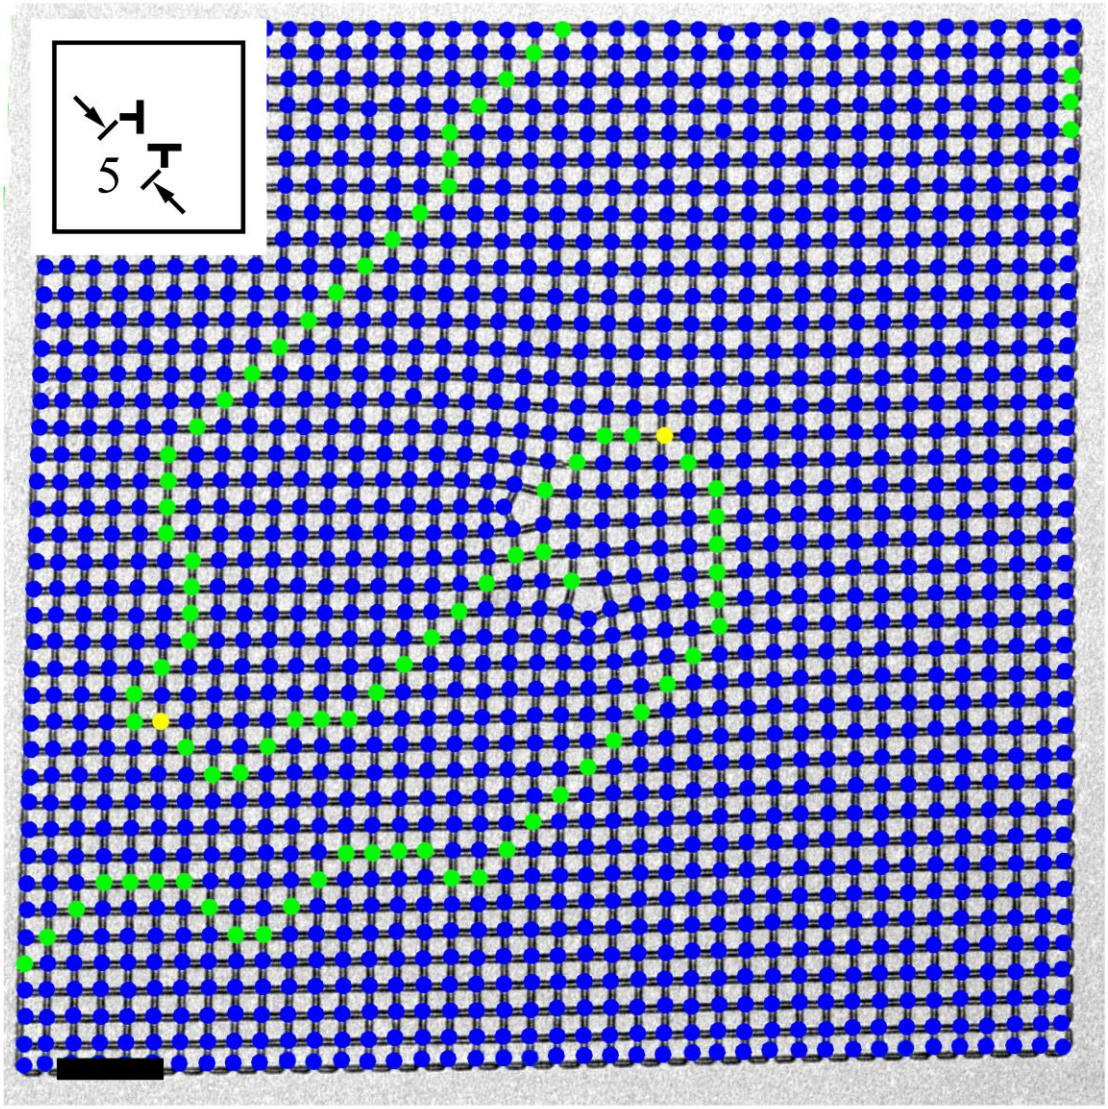

Figure 4.47 **Vertex map of a crystal with two topological defects showing meandering domain walls.** Lorentz contrast TEM image with vertex types overlaid on a *R*5 geometry crystal. A meandering vertex chain originates from each of the dislocation points and both take paths much longer than the shortest possible chain to the edges. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

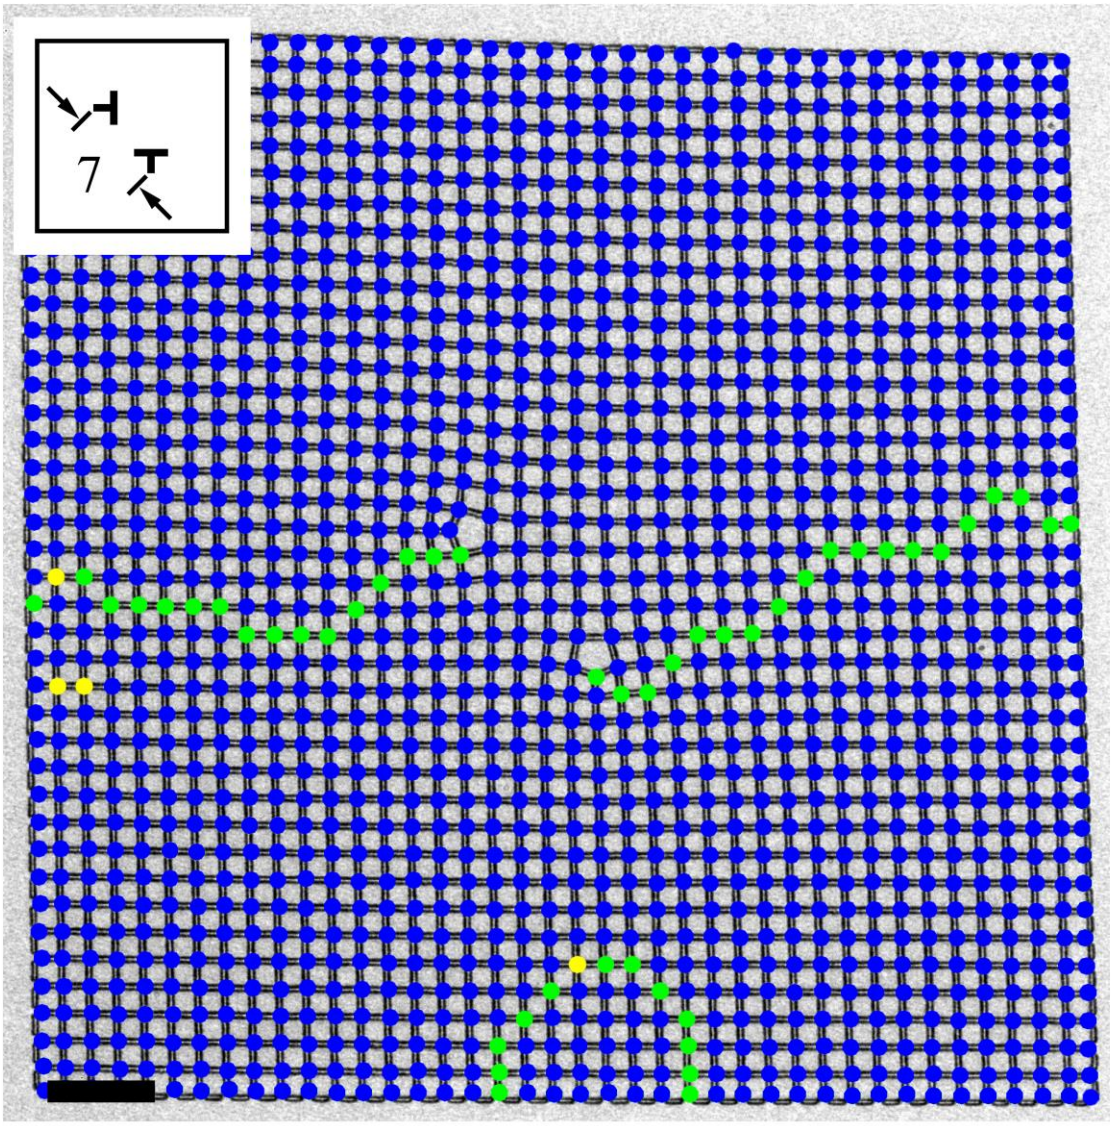

Figure 4.48 **Vertex map of a crystal with two topological defects showing relatively straight domain walls.** Lorentz contrast TEM image with vertex types overlaid on a *R7* geometry crystal. A fairly straight vertex chain originates from each of the dislocation points and both take nearly the shortest paths to the edges of the crystal. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

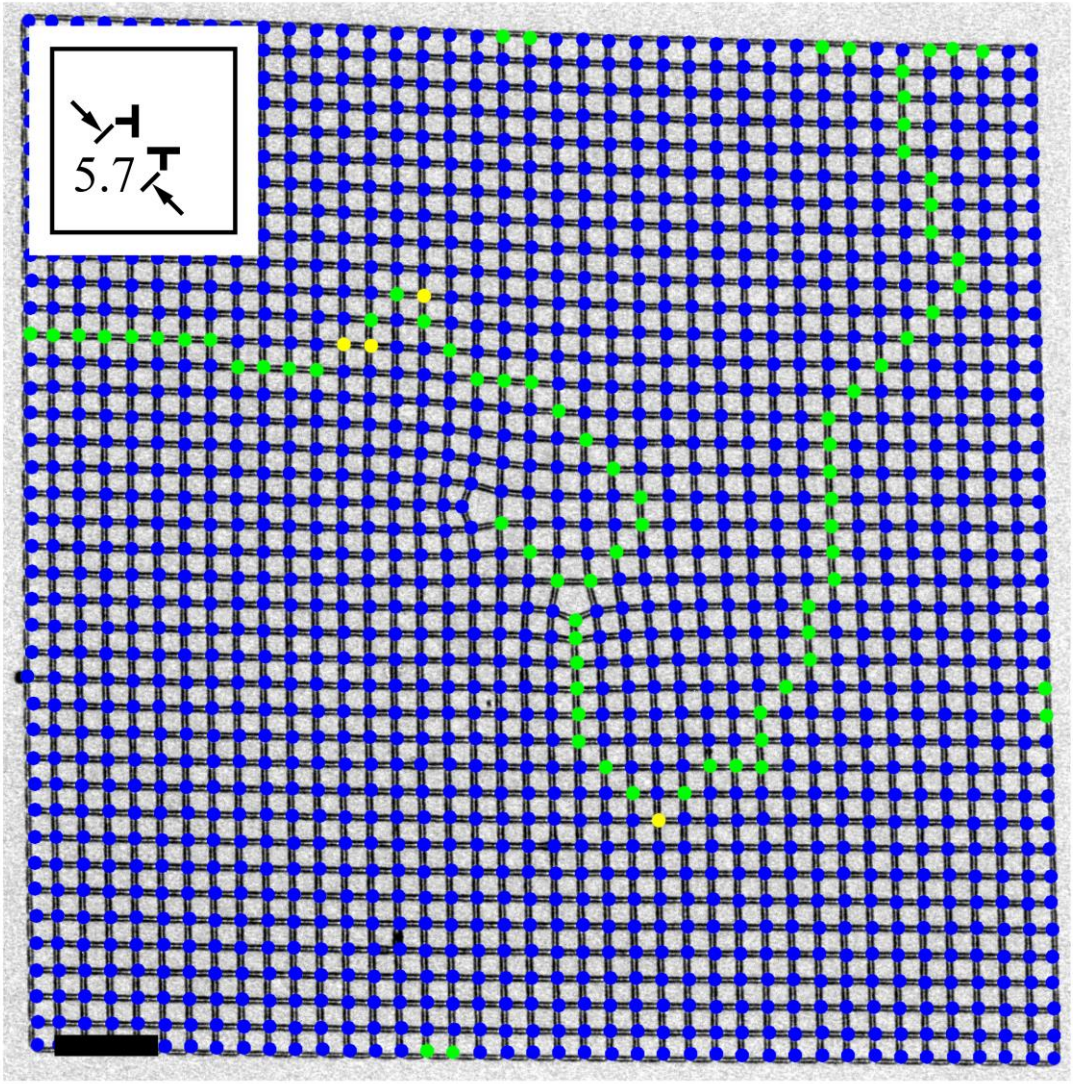

Figure 4.49 **Vertex map of a crystal with two topological defects showing unique domain wall patterns.** Lorentz contrast TEM image with vertex types overlaid on a *R*5.7 geometry crystal. A domain wall originates from one defect and connects to the other defect which also has two more domain walls originating from it and extending to the edges of the crystal. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

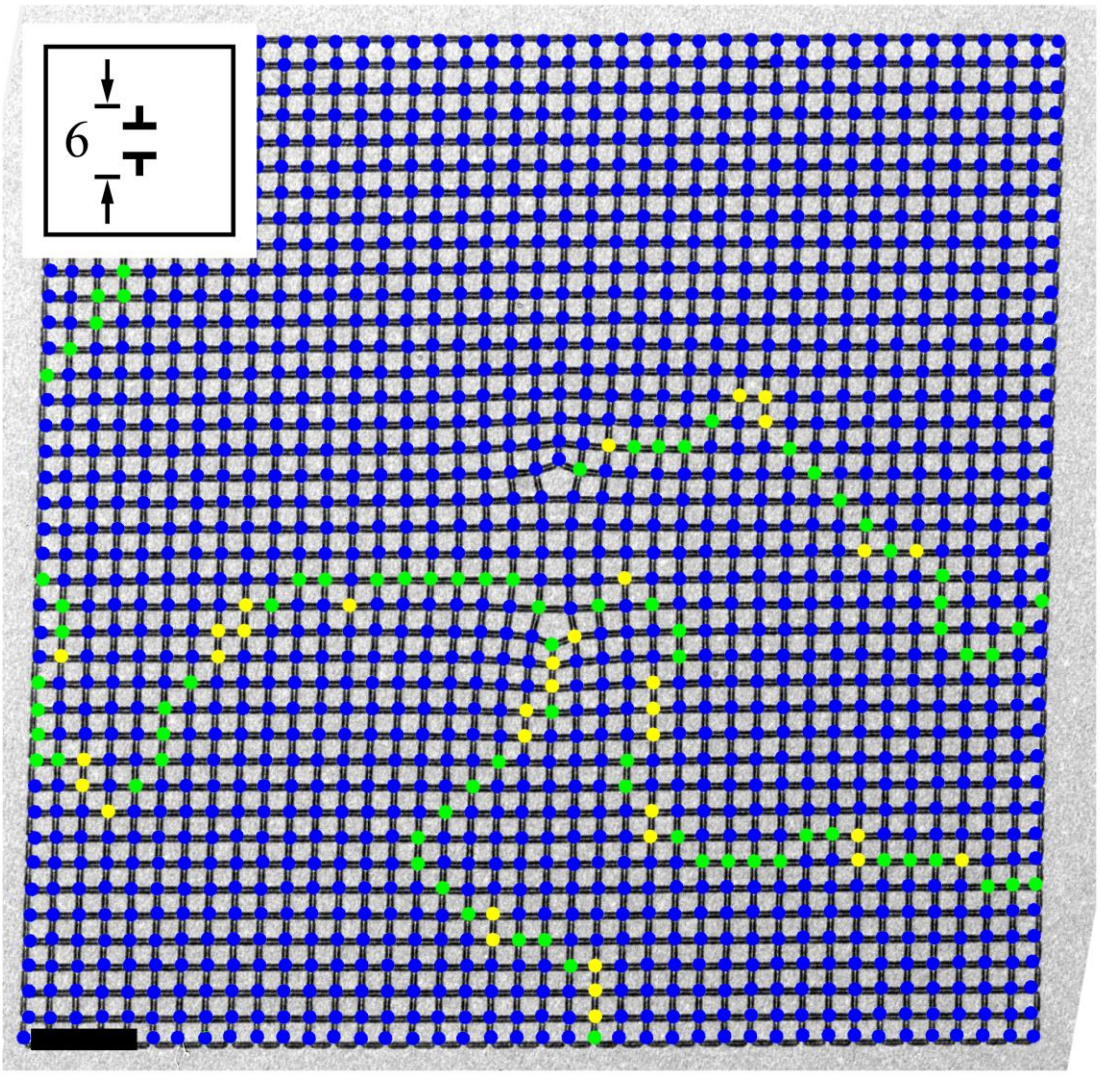

Figure 4.50 **Vertex map of a crystal with two topological defects showing unique domain wall patterns.** Lorentz contrast TEM image with vertex types overlaid on an *A*6 geometry crystal. A domain wall originates from one defect and extends to the edge of the crystal. The other defect has three domain walls that originate from it and extend to the edges of the crystal. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bar is 2 μm.

Our final observations in this section relate to the net magnetization directions of the domain walls. Both type II and type III vertices have a net magnetization associated with them. For these final images, we introduce a new notation that indicates the direction of the magnetization, shown in Figure 4.51. Type I vertices are the same blue circles and

type II and III vertices have the same colors, but the higher energy vertices are represented as arrows in the directions of the net magnetization instead of just circles. Images of two crystals with this notation are given in Figures 4.52 and 4.53. Each figure shows a large crystal with some regions enlarged that contain interesting domain walls. We see that type II vertices always have a magnetization that points in the same general direction as their neighbors. Whole chains of type II vertices all tend to follow the same direction. If the domain wall extends in a <11> direction, as in part of Figure 4.52(d), the type II vertices are all in exactly the same direction. If the wall moves along a straighter path, as in Figure 4.53(b), the neighboring vertices just have a component of the magnetization in the same direction. Figure 4.52(c) also has a nice chain of type II vertices that forms a vortex arrangement. At the type III vertices, the direction of the type II chain switches. This can be seen in Figure 4.52(b) and (d) and in Figure 4.53(c). The type III vertices are required to switch the direction of the type II chains. Overall, in this section we have seen a plethora of data and behaviors of topologically frustrated ASI crystals. In the next section, we utilize kinetic Monte Carlo simulations to model our experiments and gain a better understanding of our system.

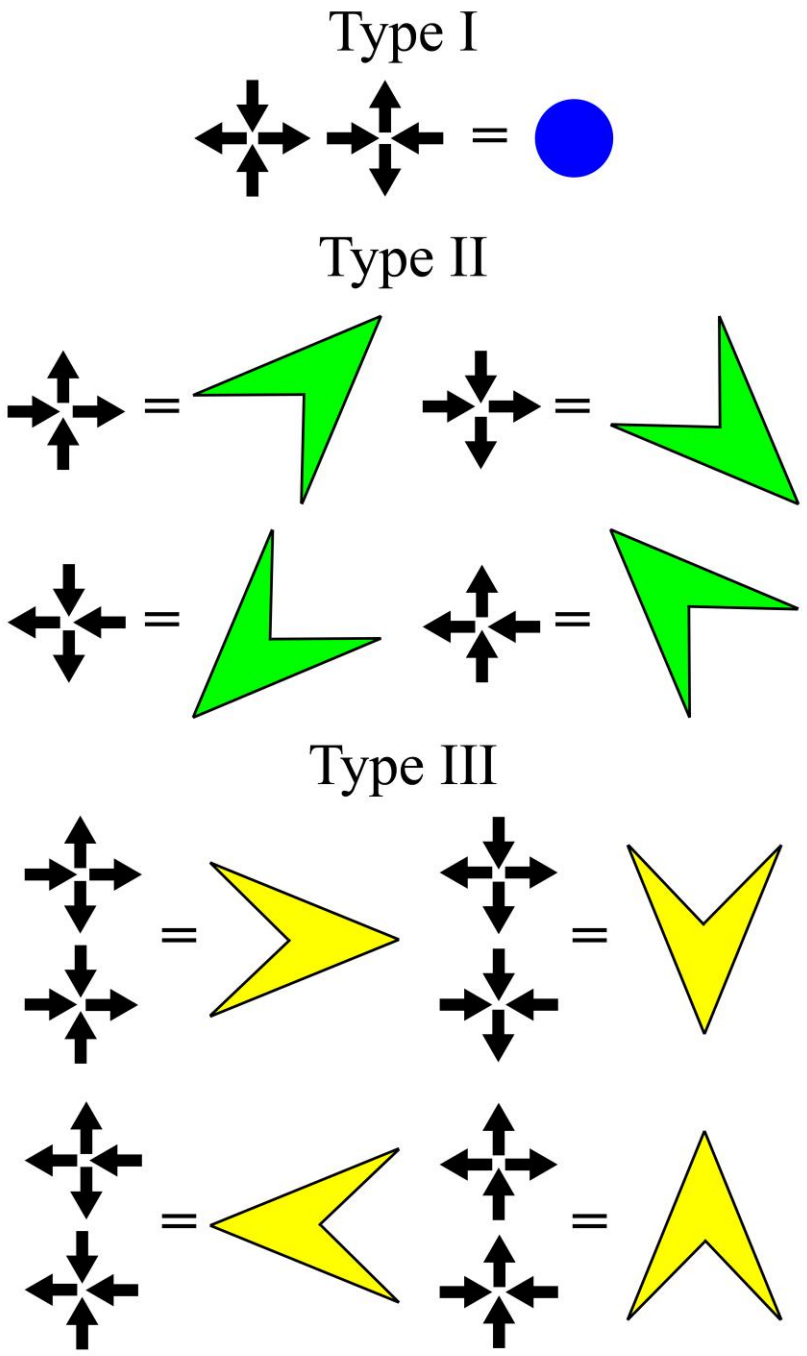

Figure 4.51 **Magnetization direction vertex notation**. Notation definitions for different vertex types used in Figures 4.52 and 4.53. Type II and type III vertices have a net magnetic moment for the vertex. This notation uses color to indicate the type of vertex (same scheme as throughout the dissertation) and an arrowhead to indicate the direction of the net magnetization of the element.

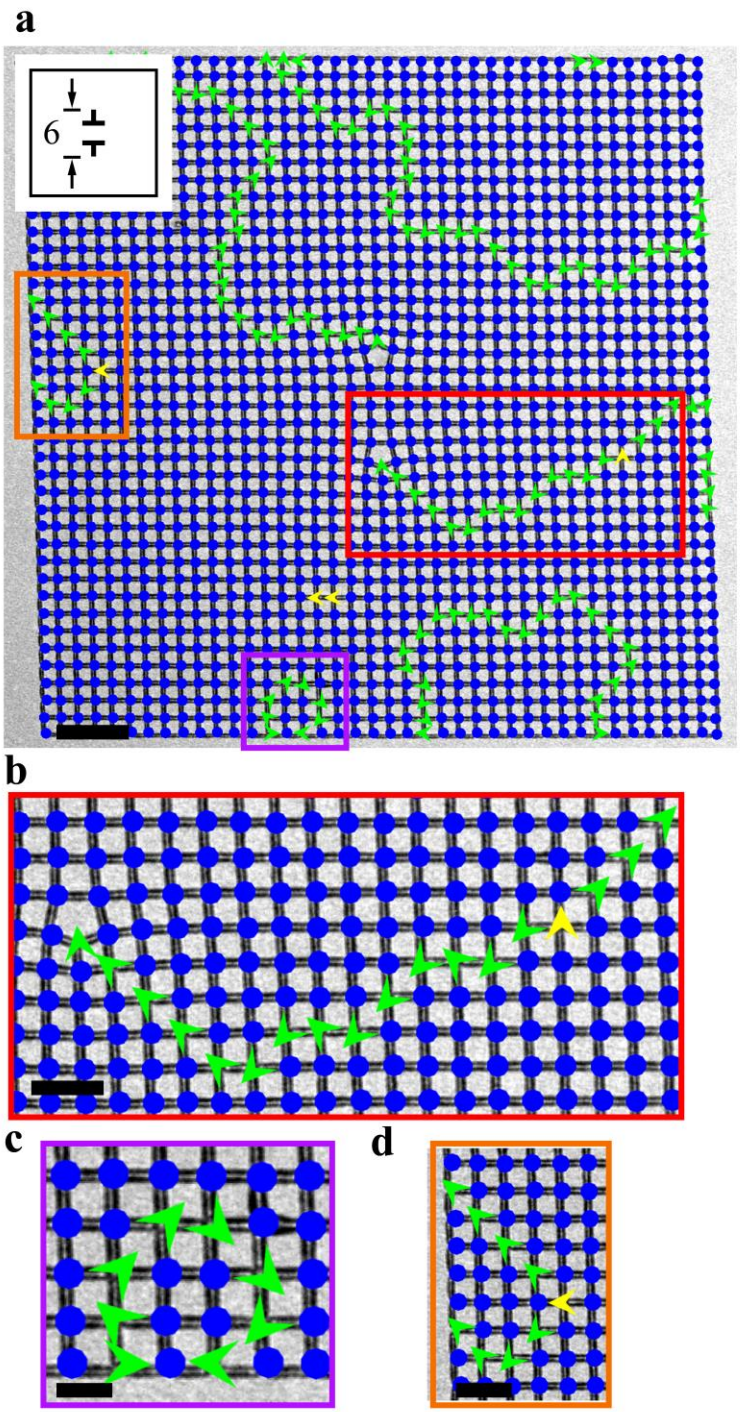

Figure 4.52 **Magnetization direction vertex map**. Lorentz TEM image of an *A*6 geometry crystal with magnetization direction vertex types overlaid. The full crystal is shown in (**a**) with regions highlighted that correspond to the magnified sections in (**b**)-(**d**). The direction of the net magnetic moments of type II vertices in a domain wall align with each other and can switch direction at a type III vertex. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bars are (**a**) 2 μm, (**b**) 1 μm, (**c**) 500 nm, and (**d**) 1 μm.

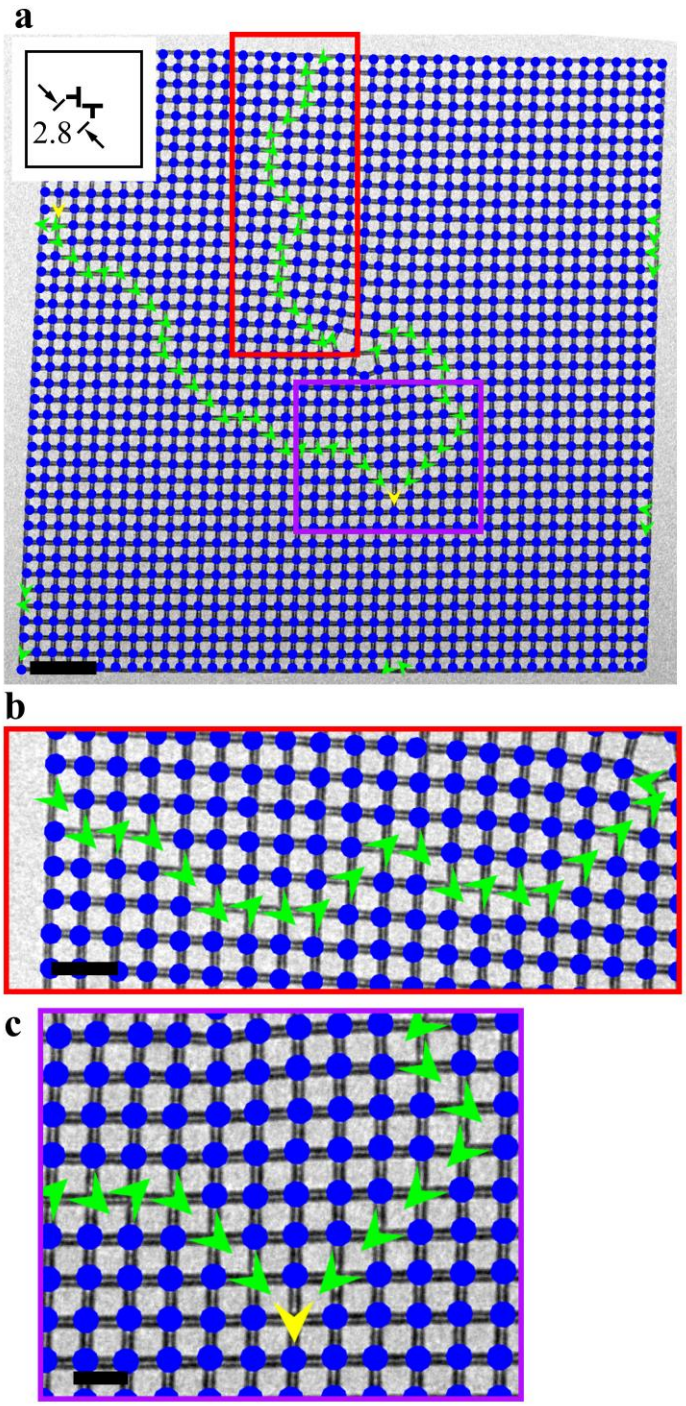

Figure 4.53 **Magnetization direction vertex map**. Lorentz TEM image of a *R*2.8 geometry crystal with magnetization direction vertex types overlaid. The full crystal is shown in (**a**) with regions highlighted that correspond to the magnified sections in (**b**) and (**c**). This is the same crystal and configuration found in Figure 4.25(a). The direction of the net magnetic moments of type II vertices in a domain wall align with each other and can switch direction at a type III vertex. Inset shows the location of the dislocation points, their directions, and the distance between them. Scale bars are (**a**) 2 μm, (**b**) 1 μm, and (**c**) 500 nm.

4.3 Kinetic Monte Carlo simulations

In this section we present results on kinetic Monte Carlo simulations of our topological defect samples. We have already discussed the numerical methods at length in Chapters 2 and 3. For these simulations, we use dipolar coupling between neighboring spins, given by the formula in Equation 1.11, and similar to the canonical square geometry simulations in Section 2.5.2. We use the same procedure to define our simulation lattices as we use for the lithography patterns. The simulations have the same number of spins and same boundary conditions as the experimental samples. We include the same type of disorder as previously discussed in the form of a spread in widths of the individual elements. We found that a small amount of disorder, on the order of 1-10 Å, was able to reproduce our experimental results extremely well. We find remarkably similar behavior between our simulations and TEM data, where the dislocation points always terminate a domain wall, but they are surrounded by large domains of ground-state order. Monte Carlo results also show domain walls separating spontaneous ordinary domains and occasionally connecting two dislocations, consistent with our experimental observations. We show the steady state behavior of some of the simulations in Figure 4.54. Here, as in the square Monte Carlo work discussed in Section 2.5.2, a given amount of disorder can generate a range of domain patterns. The frames in Figure 4.54 tended to have a relatively high number of domain walls not required by topology, but these are mainly just to illustrate the variety of patterns we see. We note that in most of the frames, the domain walls do not take the shortest path to the edge of the crystal. Two-dislocation geometry simulations show domain walls both connecting (Figure 4.54(c)) and not connecting (Figures 4.54(b) and

(d)) the dislocations. Figure 4.54(d) is relatively well ordered aside from the required domain walls.

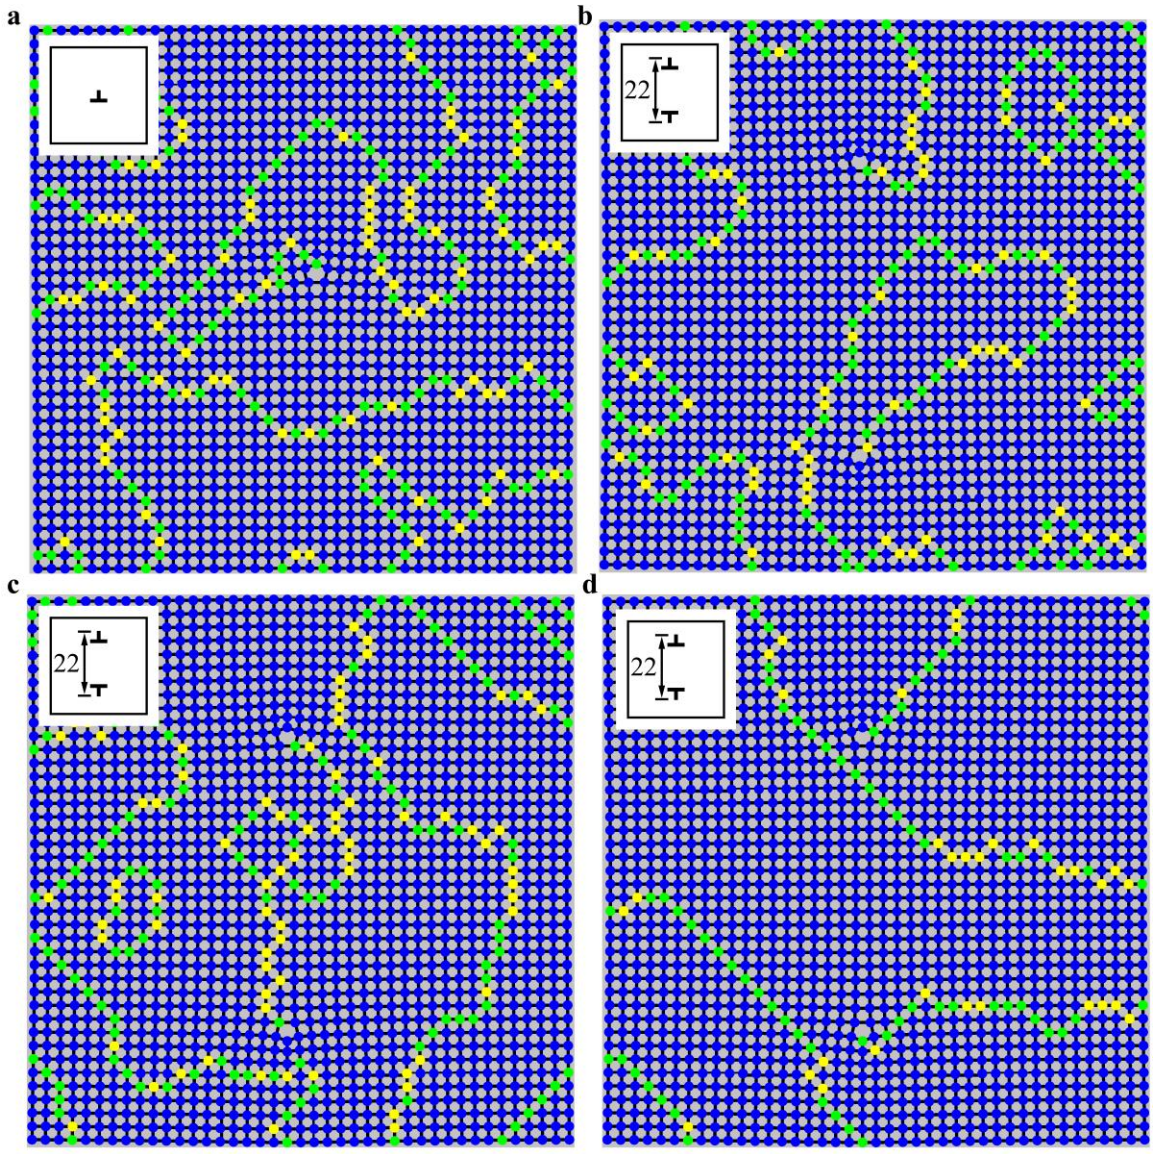

Figure 4.54 **Monte Carlo simulations.** Final frames of kinetic Monte Carlo simulations of (**a**) 1*M* and (**b**)-(**d**) *A*22 geometry lattices. The domain wall patterns observed are remarkably similar to experimental data in Section 4.2. The spread in widths used in these runs are (**a**) 5 Å, (**b**) 6 Å, (**c**) 6 Å, and (**d**) 6 Å. Insets show the location of the dislocation points, their directions, and the distance between them.

Figure 4.55 presents the evolution of the simulation shown in Figure 4.54(a). The lattice starts in a randomized state and achieves most of its ordering relatively quickly, first eliminating the highest energy type IV vertices and progressing to mainly type I ground-state order. Once the sample has reached a ground-state-domain-dominated phase, the domain walls slowly fluctuate and move around, often shrinking some of the smaller domains. In the domain wall motion, we note fluctuations of type III vertices very similar to those seen in the evolution of our canonical square geometry Monte Carlo shown in Figure 2.44, as well as previous simulations from other groups [119]. In Figure 4.55, we can see two domain walls interacting in frames (h)-(j). There are two walls in the bottom right corner of the lattice that meet and then separate into new domains. We can imagine this kind of domain wall motion is common when the crystals are thermally active and finding ordered states. Domains tend to shrink in size, favoring more ordered states, and smaller domains converge to form larger regions of ground-state order. These Monte Carlo results exactly reproduce the behavior we see in our experimental samples. Even with no disorder, we still see the topologically required domain walls originating from the dislocations. Movies of the evolution of some of these simulations as well as a movie of the lattice relaxation process can be found at dx.doi.org/10.13016/M25H71.

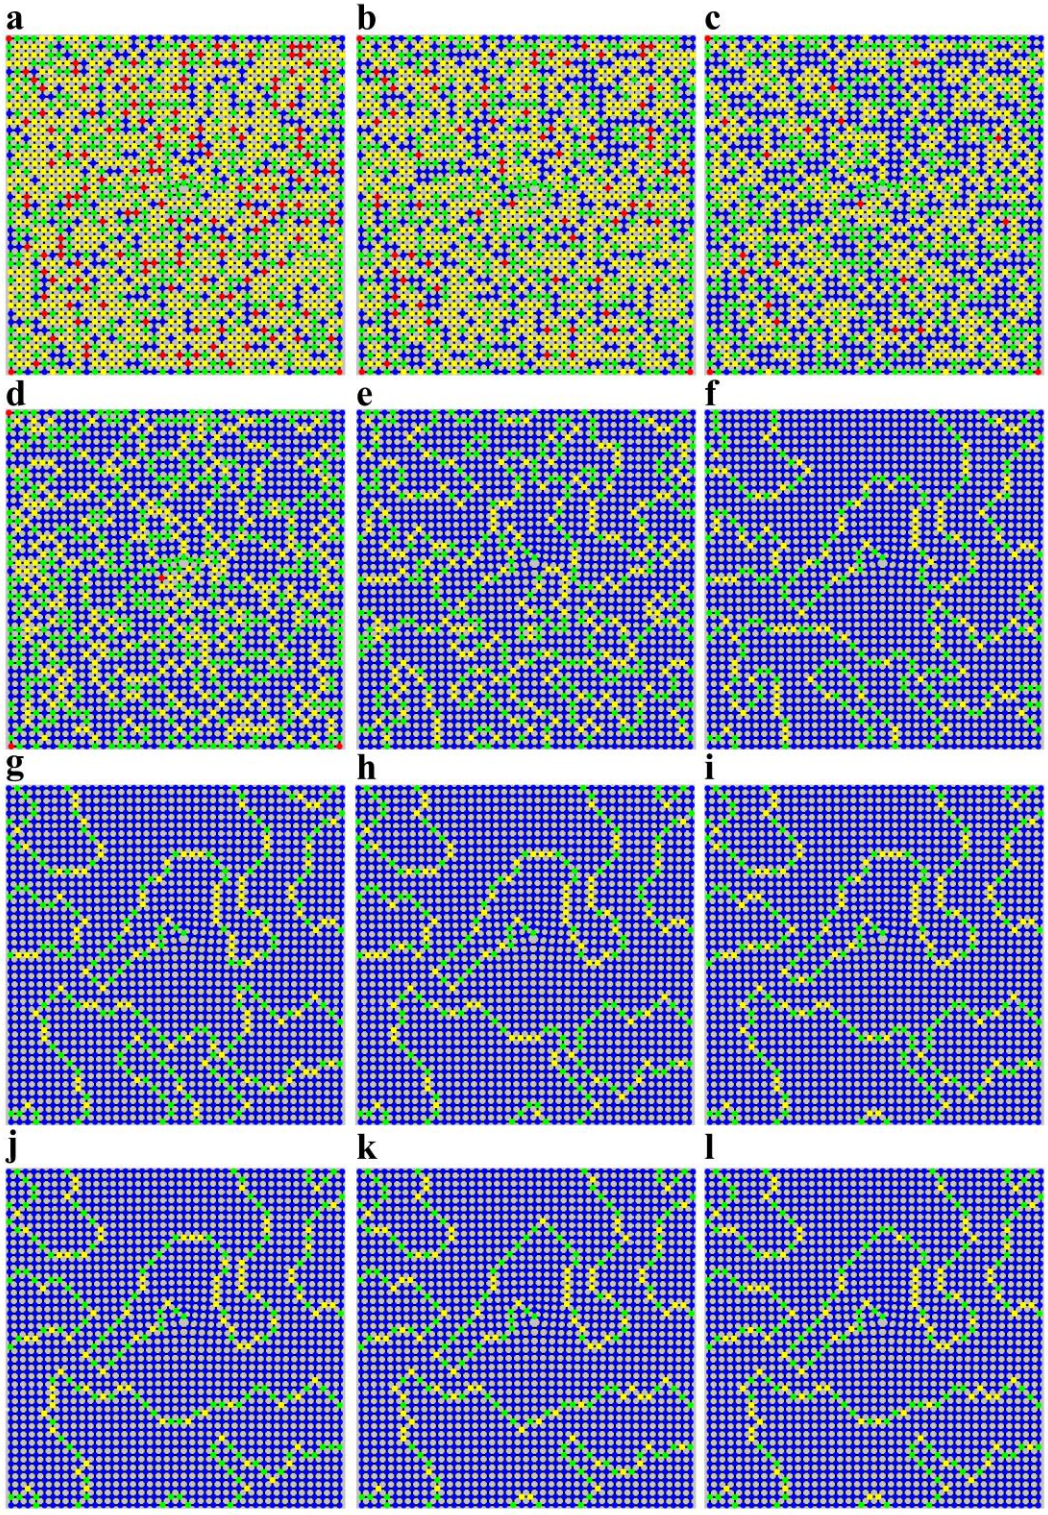

Figure 4.55 **Evolution of a topological defect kinetic Monte Carlo simulation**. Snapshots of the state of the same run as in Figure 4.54(a) after (**a**) 50, (**b**) 250, (**c**) 500, (**d**) 750, (**e**) 1250, (**f**) 2500, (**g**) 5000, (**h**) 15000, (**i**) 17500, (**j**) 20000, (**k**) 27500, and (**l**) 35000 Monte Carlo steps.

4.4 Conclusions and future work

In this chapter, we introduce controlled topological defects into canonical square geometry ASI samples and observe the resulting behavior upon annealing. We find the presence of a topological defect results in extended frustration within the system with a chain of higher energy vertices required to originate from each defect and extend to an edge of the finite lattice or to another defect among large domains of ground-state order. We develop a theoretical model for topological frustration in a variety of systems with nonzero order parameters and discuss our results in the context of unanswered questions about the pyrochlore spin ices. We provide the first experimental demonstration of dislocations nucleating domain walls with information at the single spin level. We also simulate our system and find excellent agreement between our experimental and numerical results. This work opens the door for further investigations into defects in both ASI systems and 3D crystalline materials.

In topologically frustrated ASI, there is ample more work to be done with this system. We can investigate more geometries with two dislocations to try to better understand how the dislocations behave individually and how they interact. We can fabricate larger crystals and observe how far the frustration extends into the system. We can compute the energy of different domain wall configurations to further investigate the range of behaviors we observe. We can fabricate more samples and take more data to get better statistics on the emergence of multiple dislocations not interacting with each other. We can fabricate crystals with one- or two-fold edges and see how that affects topological and non-topological domain wall and domain formation. There has been recent work on imaging thermally active magnetic configurations in real time and this would be very interesting to

compare our experimental defect crystals to our Monte Carlo movies and the evolution of our simulations [81,82,85]. We could also include more defects in the crystals and observe the effects of a higher density of dislocations in the system. The complexity of the system will grow quickly with the number of dislocations, so it is very important to have a comprehensive understanding of one- and two-defect samples before fabricating lattices with more.

The ideas presented here should be expanded on to examine the effects of other types of defects as well as investigate the behavior of defects in other ordered systems with unprecedented information at the atomic level. The work in this chapter is new and thus opens many routes for future studies. A potentially fruitful path would be to look for signatures of topological frustration in 3D crystalline materials. The pyrochlore spin ices are an obvious starting point with the aim of resolving the longstanding and ongoing debates about the ability of the system to find a ground state or whether this is even possible at all. Beyond the pyrochlores, topological frustration could have far reaching implications in a wide range of systems with nonzero order parameters. Dislocations have typically not been given much theoretical treatment, so it would be interesting to examine their effects in more detail. Moreover, our results should demand that more attention be paid to these defects as we have show they cannot and should not be ignored, for they can completely alter the properties of a system. Overall, we have expanded the field of topological frustration and presented data with unprecedented detail. ASI has proven again to be a highly tunable and useful system and will undoubtedly aid in future investigations into the properties of defects in ordered systems.

Glossary

List of acronyms

1FE – One-Fold Edge

2FE – Two-Fold Edge

3FE – Three-Fold Edge

AF – Antiferromagnetic

ASI – Artificial Spin Ice

CNAM – Center for Nanophysics and Advanced Materials

D_z – Demagnetizing factor

EBL – Electron Beam Lithography

EDS – Energy Dispersive x-ray Spectroscopy

EELS – Electron Energy Loss Spectroscopy

FCC – Face Centered Cubic

FIB – Focused Ion Beam

FM – Ferromagnetic

H_c – Coercive field

IL – Intermediate Lens

IPA – Isopropanol

LLG – Landau-Lifshitz-Gilbert

LTEM – Lorentz Transmission Electron Microscopy

MDG – Melko, den Hertog, and Gingras

MFM – Magnetic Force Microscopy

MIBK – Methyl Isobutyl Ketone

MOKE – Magneto-Optical Kerr Effect

M_r – Remnant magnetization

OOMMF – Object Oriented MicroMagnetic Framework

PEEM – Photoemission Electron Microscopy

PMMA – Poly(Methyl Methacrylate)

RKKY – Ruderman-Kittel-Kasuya-Yosida

SEM – Scanning Electron Microscopy

SQUID – Superconducting QUantum Interference Device

T_b – Blocking Temperature

T_C – Curie Temperature

TEM – Transmission Electron Microscopy

T_N – Néel Temperature

UHV – Ultra-High Vacuum

UMD – University of Maryland

WDS – Wavelength Dispersive x-ray Spectroscopy

XMCD – X-ray Magnetic Circular Dichroism

References

[1] W. Nernst, *The New Heat Theorem* (Methuen and Company, 1926).

[2] L. Pauling, "The Structure and Entropy of Ice and of Other Crystals with Some Randomness of Atomic Arrangement," *J. Am. Chem. Soc.* **57**, 2680 (1935).

[3] W. F. Giauque and M. F. Ashley, "Molecular Rotation in Ice at 10°K. Free Energy of Formation and Entropy of Water," *Phys. Rev.* **43**, 81 (1933).

[4] A. R. Gordon, "The Calculation of Thermodynamic Quantities from Spectroscopic Data for Polyatomic Molecules; the Free Energy, Entropy and Heat Capacity of Steam," *J. Chem. Phys.* **2**, 65 (1934).

[5] F. Simon, "Handbuch der Physik" (1926).

[6] W. F. Giauque and J. W. Stout, "The Entropy of Water and the Third Law of Thermodynamics. The Heat Capacity of Ice from 15 to 273°K," *J. Am. Chem. Soc.* **58**, 1144 (1936).

[7] H. P. Mulholland and R. H. Fowler, "An Asymptotic Expansion for," *Math. Proc. Cambridge Philos. Soc.* **24**, 280 (1928).

[8] G. B. B. M. Sutherland, "Rotational Specific Heat and Rotational Entropy of Simple Gases at Moderate Temperatures," *Math. Proc. Cambridge Philos. Soc.* **26**, 402 (1930).

[9] A. R. Gordon and C. Barnes, "Evaluation of the Series Which Arise in the Calculation of Thermodynamic Quantities From Spectroscopic Data," *J. Chem. Phys.* **1**, 297 (1933).

[10] V. F. Petrenko and R. W. Whitworth, *Physics of Ice* (Oxford University Press, New York, 1999).

[11] J. D. Bernal and R. H. Fowler, "A Theory of Water and Ionic Solution, with Particular Reference to Hydrogen and Hydroxyl Ions," *J. Chem. Phys.* **1**, 515 (1933).

[12] J. F. Nagle, "Lattice Statistics of Hydrogen Bonded Crystals. I. The Residual Entropy of Ice," *J. Math. Phys.* **7**, 1484 (1966).

[13] A. P. Ramirez, A. Hayashi, R. J. Cava, R. Siddharthan, and B. S. Shastry, "Zero-point entropy in 'spin ice,'" *Nature* **399**, 333 (1999).

[14] H. T. Diep, *Frustrated spin systems* (World Scientific, New Jersey, 2005).

[15] S. Franz, *Statistical Mechanics* (Springer, 2002).

[16] H. Ibach and H. Luth, *Solid-State Physics*, 3rd ed. (Springer-Verlag Berlin Heidelberg, New York, 2003).

[17] C. Kittel, *Introduction to Solid State Physics*, 8th ed. (John Wiley & Sons, 2005).

[18] M. V Berry and A. K. Geim, "Of flying frogs and levitrons," *Eur. J. Phys.* **18**, 307 (1997).

[19] W. Meissner and R. Ochsenfeld, "Ein neuer Effekt bei Eintritt der Supraleitfahigkeit," *Naturwissenschaften* **21**, 787 (1933).

[20] M. Tinkham, *Introduction to Superconductivity*, 2nd ed. (Dover, Mineola, 1996).

[21] N. W. Ashcroft and D. N. Mermin, *Solid State Physics* (Holt, Rinehart, and Winston, New York, 1976).

[22] J. D. Jackson, *Classical Electrodynamics*, 3rd ed. (John Wiley & Sons, New York, 1999).

[23] S. Ladak, S. K. Walton, K. Zeissler, T. Tyliszczak, D. E. Read, W. R. Branford, and L. F. Cohen, "Disorder-independent control of magnetic monopole defect population in artificial spin-ice honeycombs," *New J. Phys.* **14**, 045010 (2012).

[24] M. Prutton, *Thin Ferromagnetic Films* (Butterworth & Co., Washington, 1964).

[25] J. A. Osborn, "Demagnetizing Factors of the General Ellipsoid," *Phys. Rev.* **67**, 351 (1945).

[26] E. Ising, "Beitrag zur Theorie des Ferromagnetismus," *Zeitschrift für Phys.* **31**, 253 (1925).

[27] M. J. Harris, S. T. Bramwell, D. F. McMorrow, T. Zeiske, and K. W. Godfrey, "Geometrical Frustration in the Ferromagnetic Pyrochlore $Ho_2Ti_2O_7$," *Phys. Rev. Lett.* **79**, 2554 (1997).

[28] S. F. Edwards and P. W. Anderson, "Theory of spin glasses," *J. Phys. F Met. Phys.* **5**, 965 (1975).

[29] K. Binder and A. P. Young, "Spin glasses: Experimental facts, theoretical concepts, and open questions," *Rev. Mod. Phys.* **58**, 801 (1986).

[30] M. A. Ruderman and C. Kittel, "Indirect Exchange Coupling of Nuclear Magnetic Moments by Conduction Electrons," *Phys. Rev.* **96**, 99 (1954).

[31] K. Yosida, "Magnetic Properties of Cu-Mn Alloys," *Phys. Rev.* **106**, 893 (1957).

[32] T. Kasuya, "A Theory of Metallic Ferro- and Antiferromagnetism on Zener's Model," *Prog. Theor. Phys.* **16**, 45 (1956).

[33] C. A. M. Mulder, A. J. van Duyneveldt, and J. A. Mydosh, "Susceptibility of the CuMn spin-glass: Frequency and field dependences," *Phys. Rev. B* **23**, 1384 (1981).

[34] H. Maletta and W. Felsch, "Insulating spin-glass system $Eu_xSr_{1-x}S$," *Phys. Rev. B* **20**, 1245 (1979).

[35] J. S. Gardner, S. R. Dunsiger, B. D. Gaulin, M. J. P. Gingras, J. E. Greedan, R. F. Kiefl, M. D. Lumsden, W. A. MacFarlane, N. P. Raju, J. E. Sonier, I. Swainson, and Z. Tun, "Cooperative Paramagnetism in the Geometrically Frustrated Pyrochlore Antiferromagnet $Tb_2Ti_2O_7$," *Phys. Rev. Lett.* **82**, 1012 (1999).

[36] I. Mirebeau, I. N. Goncharenko, P. Cadavez-Peres, S. T. Bramwell, M. J. P. Gingras, and J. S. Gardner, "Pressure-induced crystallization of a spin liquid," *Nature* **420**, 54 (2002).

[37] L. Balents, "Spin liquids in frustrated magnets," *Nature* **464**, 199 (2010).

[38] S. Sachdev, "Quantum magnetism and criticality," *Nat. Phys.* **4**, 173 (2008).

[39] Y. Tajima, T. Matsuo, and H. Suga, "Phase transition in KOH-doped hexagonal ice," *Nature* **299**, 810 (1982).

[40] Y. K. Tsui, C. A. Burns, J. Snyder, and P. Schiffer, "Magnetic Field Induced Transitions from Spin Glass to Liquid to Long Range Order in a 3D Geometrically Frustrated Magnet," *Phys. Rev. Lett.* **82**, 3532 (1999).

[41] S.-H. Lee, C. Broholm, W. Ratcliff, G. Gasparovic, Q. Huang, T. H. Kim, and S.-W. Cheong, "Emergent excitations in a geometrically frustrated magnet," *Nature* **418**, 856 (2002).

[42] B. Martínez, X. Obradors, L. Balcells, A. Rouanet, and C. Monty, "Low Temperature Surface Spin-Glass Transition in γ-Fe_2O_3 Nanoparticles," *Phys. Rev. Lett.* **80**, 181 (1998).

[43] S. T. Bramwell and M. J. P. Gingras, "Spin Ice State in Frustrated Magnetic Pyrochlore Materials," *Science* **294**, 1495 (2001).

[44] H. Kadowaki, Y. Ishii, K. Matsuhira, and Y. Hinatsu, "Neutron scattering study of dipolar spin ice in $Ho_2Sn_2O_7$: Frustrated pyrochlore magnet," *Phys. Rev. B* **65**, 144421 (2002).

[45] J. S. Gardner, M. J. P. Gingras, and J. E. Greedan, "Magnetic pyrochlore oxides," *Rev. Mod. Phys.* **82**, 53 (2010).

[46] S. T. Bramwell, M. J. Harris, B. C. den Hertog, M. J. P. Gingras, J. S. Gardner, D. F. McMorrow, A. R. Wildes, A. L. Cornelius, J. D. M. Champion, R. G. Melko, and T. Fennell, "Spin Correlations in $Ho_2Ti_2O_7$: A Dipolar Spin Ice System," *Phys. Rev. Lett.* **87**, 047205 (2001).

[47] S. T. Bramwell, M. J. P. Gingras, and P. C. W. Holdsworth, "Spin Ice," in *Frustrated Spin Systems* (World Scientific, New Jersey, 2005).

[48] J. Rodríguez-Carvajal, "Recent advances in magnetic structure determination by neutron powder diffraction," *Phys. B Condens. Matter* **192**, 55 (1993).

[49] S. K. Sinha, E. B. Sirota, S. Garoff, and H. B. Stanley, "X-ray and neutron scattering from rough surfaces," *Phys. Rev. B* **38**, 2297 (1988).

[50] W. Hauser and H. Feshbach, "The Inelastic Scattering of Neutrons," *Phys. Rev.* **87**, 366 (1952).

[51] R. Melko, B. den Hertog, and M. Gingras, "Long-Range Order at Low Temperatures in Dipolar Spin Ice," *Phys. Rev. Lett.* **87**, 067203 (2001).

[52] R. G. Melko and M. J. P. Gingras, "Monte Carlo studies of the dipolar spin ice model," *J. Phys. Condens. Matter* **16**, 1277 (2004).

[53] B. Klemke, M. Meissner, P. Strehlow, K. Kiefer, S. A. Grigera, and D. A. Tennant, "Thermal Relaxation and Heat Transport in the Spin Ice Material $Dy_2Ti_2O_7$," *J. Low Temp. Phys.* **163**, 345 (2011).

[54] R. Higashinaka, H. Fukazawa, D. Yanagishima, and Y. Maeno, "Specific heat of $Dy_2Ti_2O_7$ in magnetic fields: comparison between single-crystalline and polycrystalline data," *J. Phys. Chem. Solids* **63**, 1043 (2002).

[55] 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 $Dy_2Ti_2O_7$," *Science* **326**, 411 (2009).

[56] H. Kadowaki, N. Doi, Y. Aoki, Y. Tabata, T. J. Sato, J. W. Lynn, K. Matsuhira, and Z. Hiroi, "Observation of Magnetic Monopoles in Spin Ice," *J. Phys. Soc. Japan* **78**, 103706 (2009).

[57] D. Pomaranski, L. R. Yaraskavitch, S. Meng, K. A. Ross, H. M. L. Noad, H. A. Dabkowska, B. D. Gaulin, and J. B. Kycia, "Absence of Pauling's residual entropy in thermally equilibrated $Dy_2Ti_2O_7$," *Nat. Phys.* **9**, 353 (2013).

[58] L. Bovo, X. Moya, D. Prabhakaran, Y.-A. Soh, A. T. Boothroyd, N. D. Mathur, G. Aeppli, and S. T. Bramwell, "Restoration of the third law in spin ice thin films," *Nat. Commun.* **5**, 1 (2014).

[59] C. Castelnovo, R. Moessner, and S. L. Sondhi, "Magnetic monopoles in spin ice," *Nature* **451**, 42 (2008).

[60] L. D. C. Jaubert and P. C. W. Holdsworth, "Signature of magnetic monopole and Dirac string dynamics in spin ice," *Nat. Phys.* **5**, 258 (2009).

[61] T. Fennell, P. P. Deen, A. R. Wildes, K. Schmalzl, D. Prabhakaran, A. T. Boothroyd, R. J. Aldus, D. F. McMorrow, and S. T. Bramwell, "Magnetic Coulomb phase in the spin ice $Ho_2Ti_2O_7$," *Science* **326**, 415 (2009).

[62] J. Snyder, B. G. Ueland, J. S. Slusky, H. Karunadasa, R. J. Cava, and P. Schiffer, "Low-temperature spin freezing in the $Dy_2Ti_2O_7$ spin ice," *Phys. Rev. B* **69**, 064414 (2004).

[63] S. T. Bramwell, S. R. Giblin, S. Calder, R. Aldus, D. Prabhakaran, and T. Fennell, "Measurement of the charge and current of magnetic monopoles in spin ice," *Nature* **461**, 956 (2009).

[64] S. J. Blundell, "Monopoles, Magnetricity, and the Stray Field from Spin Ice," *Phys. Rev. Lett.* **108**, 147601 (2012).

[65] S. R. Dunsiger, A. A. Aczel, C. Arguello, H. Dabkowska, A. Dabkowski, M.-H. Du, T. Goko, B. Javanparast, T. Lin, F. L. Ning, H. M. L. Noad, D. J. Singh, T. J. Williams, Y. J. Uemura, M. J. P. Gingras, and G. M. Luke, "Spin Ice: Magnetic Excitations without Monopole Signatures Using Muon Spin Rotation," *Phys. Rev. Lett.* **107**, 207207 (2011).

[66] L. Bovo, J. A. Bloxsom, D. Prabhakaran, G. Aeppli, and S. T. Bramwell, "Brownian motion and quantum dynamics of magnetic monopoles in spin ice.," *Nat. Commun.* **4**, 1535 (2013).

[67] C. Paulsen, M. J. Jackson, E. Lhotel, B. Canals, D. Prabhakaran, K. Matsuhira, S. R. Giblin, and S. T. Bramwell, "Far-from-equilibrium monopole dynamics in spin ice," *Nat. Phys.* **10**, 135 (2014).

[68] S. R. Giblin, S. T. Bramwell, P. C. W. Holdsworth, D. Prabhakaran, and I. Terry, "Creation and measurement of long-lived magnetic monopole currents in spin ice," *Nat. Phys.* **7**, 252 (2011).

[69] P. A. M. Dirac, "Quantised Singularities in the Electromagnetic Field," *Proc. R. Soc. A Math. Phys. Eng. Sci.* **133**, 60 (1931).

[70] B. Cabrera, "First Results from a Superconductive Detector for Moving Magnetic Monopoles," *Phys. Rev. Lett.* **48**, 1378 (1982).

[71] A. Abulencia *et al*. (CDF Collaboration), "Direct Search for Dirac Magnetic Monopoles in p pbar Collisions at Sqrt(s) = 1.96 TeV," *Phys. Rev. Lett.* **96**, 201801 (2006).

[72] A. S. Goldhaber, "Resource Letter MM-1: Magnetic monopoles," *Am. J. Phys.* **58**, 429 (1990).

[73] H. Jeon and M. J. Longo, "Search for Magnetic Monopoles Trapped in Matter," *Phys. Rev. Lett.* **75**, 1443 (1995).

[74] M. W. Ray, E. Ruokokoski, S. Kandel, M. Möttönen, and D. S. Hall, "Observation of Dirac monopoles in a synthetic magnetic field," *Nature* **505**, 657 (2014).

[75] 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* **439**, 303 (2006).

[76] C. Nisoli, R. Moessner, and P. Schiffer, "Colloquium: Artificial spin ice: Designing and imaging magnetic frustration," *Rev. Mod. Phys.* **85**, 1473 (2013).

[77] J. P. Morgan, A. Stein, S. Langridge, and C. H. Marrows, "Thermal ground-state ordering and elementary excitations in artificial magnetic square ice," *Nat. Phys.* **7**, 75 (2010).

[78] 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* **500**, 553 (2013).

[79] I. Gilbert, G.-W. Chern, S. Zhang, L. O'Brien, B. Fore, C. Nisoli, and P. Schiffer, "Emergent ice rule and magnetic charge screening from vertex frustration in artificial spin ice," *Nat. Phys.* **10**, 670 (2014).

[80] 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 (2010).

[81] A. Farhan, P. M. Derlet, A. Kleibert, A. Balan, R. V. Chopdekar, M. Wyss, L. Anghinolfi, F. Nolting, and L. J. Heyderman, "Exploring hyper-cubic energy landscapes in thermally active finite artificial spin-ice systems," *Nat. Phys.* **9**, 375 (2013).

[82] A. Farhan, P. M. Derlet, A. Kleibert, A. Balan, R. V. Chopdekar, M. Wyss, J. Perron, A. Scholl, F. Nolting, and L. J. Heyderman, "Direct Observation of Thermal Relaxation in Artificial Spin Ice," *Phys. Rev. Lett.* **111**, 057204 (2013).

[83] 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).

[84] U. B. Arnalds, A. Farhan, R. V. Chopdekar, V. Kapaklis, A. Balan, E. T. Papaioannou, M. Ahlberg, F. Nolting, L. J. Heyderman, and B. Hjörvarsson, "Thermalized ground state of artificial kagome spin ice building blocks," *Appl. Phys. Lett.* **101**, 112404 (2012).

[85] I. Gilbert, Y. Lao, I. Carrasquillo, L. O'Brien, J. D. Watts, M. Manno, C. Leighton, A. Scholl, C. Nisoli, and P. Schiffer, "Emergent reduced dimensionality by vertex frustration in artificial spin ice," *Nat. Phys.* **12**, 162 (2015).

[86] K. K. Kohli, A. L. Balk, J. Li, S. Zhang, I. Gilbert, P. E. Lammert, V. H. Crespi, P. Schiffer, and N. Samarth, "Magneto-optical Kerr effect studies of square artificial spin ice," *Phys. Rev. B* **84**, 180412 (2011).

[87] V. Kapaklis, U. B. Arnalds, A. Harman-Clarke, E. T. Papaioannou, M. Karimipour, P. Korelis, A. Taroni, P. C. W. Holdsworth, S. T. Bramwell, and B. Hjörvarsson, "Melting artificial spin ice," *New J. Phys.* **14**, 035009 (2012).

[88] J. P. Morgan, C. J. Kinane, T. R. Charlton, A. Stein, C. Sánchez-Hanke, D. A. Arena, S. Langridge, and C. H. Marrows, "Magnetic hysteresis of an artificial square ice studied by in-plane Bragg x-ray resonant magnetic scattering," *AIP Adv.* **2**, 022163 (2012).

[89] J. Perron, L. Anghinolfi, B. Tudu, N. Jaouen, J.-M. Tonnerre, M. Sacchi, F. Nolting, J. Lüning, and L. J. Heyderman, "Extended reciprocal space observation of artificial spin ice with x-ray resonant magnetic scattering," *Phys. Rev. B* **88**, 214424 (2013).

[90] S. D. Pollard, V. Volkov, and Y. Zhu, "Propagation of magnetic charge monopoles and Dirac flux strings in an artificial spin-ice lattice," *Phys. Rev. B* **85**, 180402 (2012).

[91] V. V. Volkov and Y. Zhu, "Lorentz phase microscopy of magnetic materials," *Ultramicroscopy* **98**, 271 (2004).

[92] C. Phatak, A. K. Petford-Long, O. Heinonen, M. Tanase, and M. De Graef, "Nanoscale structure of the magnetic induction at monopole defects in artificial spin-ice lattices," *Phys. Rev. B* **83**, 174431 (2011).

[93] 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).

[94] S. A. Daunheimer, O. Petrova, O. Tchernyshyov, and J. Cumings, "Reducing Disorder in Artificial Kagome Ice," *Phys. Rev. Lett.* **107**, 167201 (2011).

[95] J. Drisko, S. Daunheimer, and J. Cumings, "$FePd_3$ as a material for studying thermally active artificial spin ice systems," *Phys. Rev. B* **91**, 224406 (2015).

[96] J. Snyder, J. S. Slusky, R. J. Cava, and P. Schiffer, "How 'spin ice' freezes," *Nature* **413**, 48 (2001).

[97] G. Barkema and M. Newman, "Monte Carlo simulation of ice models," *Phys. Rev. E* **57**, 1155 (1998).

[98] G.-W. Chern and O. Tchernyshyov, "Magnetic charge and ordering in kagome spin ice," *Philos. Trans. A. Math. Phys. Eng. Sci.* **370**, 5718 (2012).

[99] D. Davidović, S. Kumar, D. Reich, J. Siegel, S. Field, R. Tiberio, R. Hey, and K. Ploog, "Correlations and Disorder in Arrays of Magnetically Coupled Superconducting Rings," *Phys. Rev. Lett.* **76**, 815 (1996).

[100] R. P. Cowburn and M. E. Welland, "Room Temperature Magnetic Quantum Cellular Automata," *Science* **287**, 1466 (2000).

[101] H. Hilgenkamp, Ariando, H.-J. H. Smilde, D. H. A. Blank, G. Rijnders, H. Rogalla, J. R. Kirtley, and C. C. Tsuei, "Ordering and manipulation of the magnetic moments in large-scale superconducting pi-loop arrays," *Nature* **422**, 50 (2003).

[102] R. A. Shelby, "Experimental Verification of a Negative Index of Refraction," *Science* **292**, 77 (2001).

[103] D. R. Smith, W. J. Padilla, D. C. Vier, S. C. Nemat-Nasser, and S. Schultz, "Composite Medium with Simultaneously Negative Permeability and Permittivity," *Phys. Rev. Lett.* **84**, 4184 (2000).

[104] D. R. Smith, "Metamaterials and Negative Refractive Index," *Science* **305**, 788 (2004).

[105] J. B. Pendry, "Controlling Electromagnetic Fields," *Science* **312**, 1780 (2006).

[106] X. Ni, Z. J. Wong, M. Mrejen, Y. Wang, and X. Zhang, "An ultrathin invisibility skin cloak for visible light," *Science* **349**, 1310 (2015).

[107] W. Cai, U. K. Chettiar, A. V. Kildishev, and V. M. Shalaev, "Optical cloaking with metamaterials," *Nat. Photonics* **1**, 224 (2007).

[108] D. Schurig, J. J. Mock, B. J. Justice, S. A. Cummer, J. B. Pendry, A. F. Starr, and D. R. Smith, "Metamaterial Electromagnetic Cloak at Microwave Frequencies," *Science* **314**, 977 (2006).

[109] D. A. Gilbert, B. B. Maranville, A. L. Balk, B. J. Kirby, P. Fischer, D. T. Pierce, J. Unguris, J. A. Borchers, and K. Liu, "Realization of ground-state artificial skyrmion lattices at room temperature," *Nat. Commun.* **6**, 8462 (2015).

[110] N. Lazarides and G. P. Tsironis, "rf superconducting quantum interference device metamaterials," *Appl. Phys. Lett.* **90**, 163501 (2007).

[111] M. Trepanier, D. Zhang, O. Mukhanov, and S. M. Anlage, "Realization and Modeling of Metamaterials Made of rf Superconducting Quantum-Interference Devices," *Phys. Rev. X* **3**, 041029 (2013).

[112] G. Möller and R. Moessner, "Artificial Square Ice and Related Dipolar Nanoarrays," *Phys. Rev. Lett.* **96**, 237202 (2006).

[113] R. F. Wang, J. Li, W. McConville, C. Nisoli, X. Ke, J. W. Freeland, V. Rose, M. Grimsditch, P. Lammert, V. H. Crespi, and P. Schiffer, "Demagnetization protocols for frustrated interacting nanomagnet arrays," *J. Appl. Phys.* **101**, 09J104 (2007).

[114] X. Ke, J. Li, C. Nisoli, P. Lammert, W. McConville, R. Wang, V. Crespi, and P. Schiffer, "Energy Minimization and ac Demagnetization in a Nanomagnet Array," *Phys. Rev. Lett.* **101**, 037205 (2008).

[115] C. Nisoli, R. Wang, J. Li, W. McConville, P. Lammert, P. Schiffer, and V. Crespi, "Ground State Lost but Degeneracy Found: The Effective Thermodynamics of Artificial Spin Ice," *Phys. Rev. Lett.* **98**, 217203 (2007).

[116] C. Nisoli, J. Li, X. Ke, D. Garand, P. Schiffer, and V. H. Crespi, "Effective Temperature in an Interacting Vertex System: Theory and Experiment on Artificial Spin Ice," *Phys. Rev. Lett.* **105**, 047205 (2010).

[117] J. P. Morgan, A. Bellew, A. Stein, S. Langridge, and C. H. Marrows, "Linear field demagnetization of artificial magnetic square ice," *Front. Phys.* **1**, 1 (2013).

[118] Z. Budrikis, P. Politi, and R. L. Stamps, "Diversity Enabling Equilibration: Disorder and the Ground State in Artificial Spin Ice," *Phys. Rev. Lett.* **107**, 217204 (2011).

[119] Z. Budrikis, K. L. Livesey, J. P. Morgan, J. Akerman, A. Stein, S. Langridge, C. H. Marrows, and R. L. Stamps, "Domain dynamics and fluctuations in artificial square ice at finite temperatures," *New J. Phys.* **14**, 035014 (2012).

[120] Z. Budrikis, J. P. Morgan, J. Akerman, A. Stein, P. Politi, S. Langridge, C. H. Marrows, and R. L. Stamps, "Disorder Strength and Field-Driven Ground State Domain Formation in Artificial Spin Ice: Experiment, Simulation, and Theory," *Phys. Rev. Lett.* **109**, 037203 (2012).

[121] Z. Budrikis, P. Politi, and R. L. Stamps, "Disorder regimes and equivalence of disorder types in artificial spin ice," *J. Appl. Phys.* **111**, 07E109 (2012).

[122] C. Kittel and H. Kroemer, *Thermal Physics*, 2nd ed. (W. H. Freeman, San Francisco, 1980).

[123] P. E. Lammert, X. Ke, J. Li, C. Nisoli, D. M. Garand, V. H. Crespi, and P. Schiffer, "Direct entropy determination and application to artificial spin ice," *Nat. Phys.* **6**, 786 (2010).

[124] 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).

[125] G.-W. Chern, P. Mellado, and O. Tchernyshyov, "Two-Stage Ordering of Spins in Dipolar Spin Ice on the Kagome Lattice," *Phys. Rev. Lett.* **106**, 207202 (2011).

[126] K. Matsuhira, Z. Hiroi, T. Tayama, S. Takagi, and T. Sakakibara, "A new macroscopically degenerate ground state in the spin ice compound $Dy_2Ti_2O_7$ under a magnetic field," *J. Phys. Condens. Matter* **14**, L559 (2002).

[127] G. Möller and R. Moessner, "Magnetic multipole analysis of kagome and artificial spin-ice dipolar arrays," *Phys. Rev. B* **80**, 140409 (2009).

[128] S. Ladak, D. E. Read, G. K. Perkins, L. F. Cohen, and W. R. Branford, "Direct observation of magnetic monopole defects in an artificial spin-ice system," *Nat. Phys.* **6**, 359 (Nature Publishing Group, 2010).

[129] I. Gilbert, G.-W. Chern, B. Fore, Y. Lao, S. Zhang, C. Nisoli, and P. Schiffer, "Direct visualization of memory effects in artificial spin ice," *Phys. Rev. B* **92**, 104417 (2015).

[130] A. Libál, C. J. O. Reichhardt, and C. Reichhardt, "Creating Artificial Ice States Using Vortices in Nanostructured Superconductors," *Phys. Rev. Lett.* **102**, 237004 (2009).

[131] A. Libál, C. J. O. Reichhardt, and C. Reichhardt, "Doped colloidal artificial spin ice," *New J. Phys.* **17**, 103010 (2015).

[132] S. Daunheimer, "Magnetic Reversal of Artificial Kagome Ice" (Doctoral dissertation, University of Maryland, 2013).

[133] B. Gutenberg and C. F. Richter, "Magnitude and Energy of Earthquakes," in *Ann. Geophys.* **9** (1956).

[134] M. Kardar, "Avalanche theory in rice," *Nature* **379**, 22 (1996).

[135] H. M. Jaeger, C. Liu, and S. R. Nagel, "Relaxation at the Angle of Repose," *Phys. Rev. Lett.* **62**, 40 (1989).

[136] V. Frette, K. Christensen, A. Malthe-Sørenssen, J. Feder, T. Jøssang, and P. Meakin, "Avalanche dynamics in a pile of rice," *Nature* **379**, 49 (1996).

[137] S. Wolfram, "Statistical mechanics of cellular automata," *Rev. Mod. Phys.* **55**, 601 (1983).

[138] P. Bak, C. Tang, and K. Wiesenfeld, "Self-organized criticality: An explanation of the 1/f noise," *Phys. Rev. Lett.* **59**, 381 (1987).

[139] J. M. Porro, A. Bedoya-Pinto, A. Berger, and P. Vavassori, "Exploring thermally induced states in square artificial spin-ice arrays," *New J. Phys.* **15**, 055012 (2013).

[140] F. Montaigne, D. Lacour, I. A. Chioar, N. Rougemaille, D. Louis, S. Mc Murtry, H. Riahi, B. S. Burgos, T. O. Menteş, A. Locatelli, B. Canals, and M. Hehn, "Size distribution of magnetic charge domains in thermally activated but out-of-equilibrium artificial spin ice," *Sci. Rep.* **4**, 5702 (2014).

[141] I. A. Chioar, B. Canals, D. Lacour, M. Hehn, B. Santos Burgos, T. O. Menteş, A. Locatelli, F. Montaigne, and N. Rougemaille, "Kinetic pathways to the magnetic charge crystal in artificial dipolar spin ice," *Phys. Rev. B* **90**, 220407 (2014).

[142] S. Arrhenius, "On the reaction rate of the inversion of non-refined sugar upon souring," *Z. Phys Chem* **4**, 226 (1889).

[143] J. I. Gittleman, B. Abeles, and S. Bozowski, "Superparamagnetism and relaxation effects in granular Ni-SiO_2 and Ni-Al_2O_3 films," *Phys. Rev. B* **9**, 3891 (1974).

[144] L. Neel, "Théorie du trainage magnétique des ferromagnétiques au grains fin avec applications aux terres cuites," *Ann. Géophysique* **5**, 99 (1949).

[145] A. Aharoni, "Demagnetizing factors for rectangular ferromagnetic prisms," *J. Appl. Phys.* **83**, 3432 (1998).

[146] R. J. Wakelin and E. L. Yates, "A Study of the Order-Disorder Transformation in Iron-Nickel Alloys in the Region $FeNi_3$," *Proc. Phys. Soc. Sect. B* **66**, 221 (1953).

[147] G.-W. Chern, M. J. Morrison, and C. Nisoli, "Degeneracy and Criticality from Emergent Frustration in Artificial Spin Ice," *Phys. Rev. Lett.* **111**, 177201 (2013).

[148] M. J. Morrison, T. R. Nelson, and C. Nisoli, "Unhappy vertices in artificial spin ice: new degeneracies from vertex frustration," *New J. Phys.* **15**, 045009 (2013).

[149] G. Chern, C. Reichhardt, and C. Nisoli, "Realizing three-dimensional artificial spin ice by stacking planar nano-arrays," *Appl. Phys. Lett.* **104**, 013101 (2014).

[150] S. Zhang, J. Li, I. Gilbert, J. Bartell, M. J. Erickson, Y. Pan, P. E. Lammert, C. Nisoli, K. K. Kohli, R. Misra, V. H. Crespi, N. Samarth, C. Leighton, and P. Schiffer, "Perpendicular Magnetization and Generic Realization of the Ising Model in Artificial Spin Ice," *Phys. Rev. Lett.* **109**, 087201 (2012).

[151] V. S. Bhat, J. Sklenar, B. Farmer, J. Woods, J. T. Hastings, S. J. Lee, J. B. Ketterson, and L. E. De Long, "Controlled Magnetic Reversal in Permalloy Films Patterned into Artificial Quasicrystals," *Phys. Rev. Lett.* **111**, 077201 (2013).

[152] V. S. Bhat, B. Farmer, N. Smith, E. Teipel, J. Woods, J. Sklenar, J. B. Ketterson, J. T. Hastings, and L. E. De Long, "Non-stochastic switching and emergence of magnetic vortices in artificial quasicrystal spin ice," *Phys. C Supercond.* **503**, (2014).

[153] P. Mellado, A. Concha, and L. Mahadevan, "Macroscopic Magnetic Frustration," *Phys. Rev. Lett.* **109**, 257203 (2012).

[154] J. Trastoy, M. Malnou, C. Ulysse, R. Bernard, N. Bergeal, G. Faini, J. Lesueur, J. Briatico, and J. E. Villegas, "Freezing and thawing of artificial ice by thermal switching of geometric frustration in magnetic flux lattices," *Nat. Nanotechnol.* **9**, 710 (2014).

[155] M. L. Latimer, G. R. Berdiyorov, Z. L. Xiao, F. M. Peeters, and W. K. Kwok, "Realization of Artificial Ice Systems for Magnetic Vortices in a Superconducting MoGe Thin Film with Patterned Nanostructures," *Phys. Rev. Lett.* **111**, 067001 (2013).

[156] A. Libál, C. Reichhardt, and C. J. Olson Reichhardt, "Hysteresis and return-point memory in colloidal artificial spin ice systems," *Phys. Rev. E* **86**, 021406 (2012).

[157] G.-W. Chern, C. Reichhardt, and C. J. Olson Reichhardt, "Frustrated colloidal ordering and fully packed loops in arrays of optical traps," *Phys. Rev. E* **87**, 062305 (2013).

[158] Y. Han, Y. Shokef, A. M. Alsayed, P. Yunker, T. C. Lubensky, and A. G. Yodh, "Geometric frustration in buckled colloidal monolayers," *Nature* **456**, 898 (2008).

[159] A. Libál, C. Reichhardt, and C. J. O. Reichhardt, "Realizing Colloidal Artificial Ice on Arrays of Optical Traps," *Phys. Rev. Lett.* **97**, 228302 (2006).

[160] C. J. Olson Reichhardt, A. Libál, and C. Reichhardt, "Multi-step ordering in kagome and square artificial spin ice," *New J. Phys.* **14**, 025006 (2012).

[161] Y. Shokef, Y. Han, A. Souslov, A. G. Yodh, and T. C. Lubensky, "Buckled colloidal monolayers connect geometric frustration in soft and hard matter," *Soft Matter* **9**, 6565 (2013).

[162] A. Kussmann and K. Jessen, "Magnetische und Dilatometrische Messungen zur Umwandlungskinetik der Eisen-Palladium-Legierungen," *Zeitschrift fur Met.* **54**, 504 (1963).

[163] G. Longworth, "Temperature Dependence of the ^{57}Fe hfs in the Ordered Alloys $FePd_3$ and FePd neat the Curie Temperature," *Phys. Rev.* **172**, 572 (1968).

[164] E. Burzo and P. Vlaic, "Magnetic properties of iron-palladium solid solutions and compounds," *J. Optoelectron. Adv. Mater.* **12**, 1869 (2010).

[165] V. G. Myagkov, V. S. Zhigalov, B. A. Belyaev, L. E. Bykova, L. A. Solovyov, and G. N. Bondarenko, "Solid-state synthesis and magnetic properties of epitaxial FePd3(001) films," *J. Magn. Magn. Mater.* **324**, 1571 (2012).

[166] B. P. Burton, *Fe-Pd (Iron-Palladium), Binary Alloy Phase Diagrams*, 2nd ed., T. B. Massalski, Ed. (1990).

[167] A. C. Johnston-Peck, G. Scarel, J. Wang, G. N. Parsons, and J. B. Tracy, "Sinter-free phase conversion and scanning transmission electron microscopy of FePt nanoparticle monolayers," *Nanoscale* **3**, 4142 (2011).

[168] B. Window and G. Longworth, "Electrostatic interactions at ^{57}Fe nuclei in nickel, palladium and platinum alloys," *J. Phys. F Met. Phys.* **1**, 325 (1971).

[169] A. M. Clogston, B. T. Matthias, M. Peter, H. J. Williams, E. Corenzwit, and R. C. Sherwood, "Local Magnetic Moment Associated with an Iron Atom Dissolved in Various Transition Metal Alloys," *Phys. Rev.* **125**, 541 (1962).

[170] J. B. Staunton, "The electronic structure of magnetic transition metallic materials," *Reports Prog. Phys.* **57**, 1289 (1994).

[171] D. Adachi, D. Bonnenberg, J. J. M. Franse, R. Gersdorf, K. A. Hempel, K. Kanematse, S. Misawa, M. Shiga, and M. B. Stearns, *Landolt Börnstein Handbuch vol. 19a* (Springer Verlag, 1992).

[172] M. Matsui and K. Adachi, "Magneto-elastic properties and invar anomaly of Fe-Pd alloys," *Phys. B Condens. Matter* **161**, 53 (1990).

[173] "Silson," <http://www.silson.com/index.html?content=tem>.

[174] Sigma-Aldrich, "Physical Vapor Deposition," <https://www.sigmaaldrich.com/materials-science/material-science-products.html?TablePage=108832720>.

[175] D. B. Williams and C. B. Carter, *Transmission Electron Microscopy: A Textbook for Materials Science*, 2nd ed. (Springer Science + Business Media, LLC, New York, 2009).

[176] A. N. Broers, A. C. F. Hoole, and J. M. Ryan, "Electron beam lithography—Resolution limits," *Microelectron. Eng.* **32**, 131 (1996).

[177] C. Shawn Wu, Y. Makiuchi, and C. Chen, "High-energy Electron Beam Lithography for Nanoscale Fabrication," in *Lithography*, M. Wang, Ed. (InTech, 2010).

[178] JEOL, "Electron Beam Source for Electron Beam Deposition," <http://www.jeol.co.jp/en/science/eb.html>.

[179] JEOL, "JEM-2100 Electron Microscope Manual" (2004).

[180] Z. L. Wang, "Fundamental Theory of Transmission Electronic Microscopy," <http://www.nanoscience.gatech.edu/zlwang/research/tem.html>.

[181] Y. Qi, "Artificial Kagome Spin Ice" (Doctoral dissertation, University of Maryland, 2008).

[182] L. Reimer and H. Kohl, *Transmission Electron Microscopy: Physics of Image Formation*, 5th ed. (Springer Science + Business Media, LLC, New York, 2008).

[183] D. Levis, L. F. Cugliandolo, L. Foini, and M. Tarzia, "Thermal Phase Transitions in Artificial Spin Ice," *Phys. Rev. Lett.* **110**, 207206 (2013).

[184] I. A. Chioar, N. Rougemaille, A. Grimm, O. Fruchart, E. Wagner, M. Hehn, D. Lacour, F. Montaigne, and B. Canals, "Nonuniversality of artificial frustrated spin systems," *Phys. Rev. B* **90**, 064411 (2014).

[185] C. Jacoboni and L. Reggiani, "The Monte Carlo method for the solution of charge transport in semiconductors with applications to covalent materials," *Rev. Mod. Phys.* **55**, 645 (1983).

[186] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, "Quantum Monte Carlo simulations of solids," *Rev. Mod. Phys.* **73**, 33 (2001).

[187] W. K. Hastings, "Monte Carlo Sampling Methods Using Markov Chains and Their Applications," *Biometrika* **57**, 97 (1970).

[188] S. Chib and E. Greenberg, "Understanding the Metropolis-Hastings Algorithm," *Am. Stat.* **49**, 327 (1995).

[189] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of State Calculations by Fast Computing Machines," *J. Chem. Phys.* **21**, 1087 (1953).

[190] W. M. Young and E. W. Elcock, "Monte Carlo studies of vacancy migration in binary ordered alloys: I," *Proc. Phys. Soc.* **89**, 735 (1966).

[191] G. Chern and O. Tchernyshyov, "Magnetic charge and ordering in kagome spin ice," *ArXiv*:1109.0275 (2008).

[192] M. J. Donahue and D. G. Porter, "OOMMF User's Guide, Version 1.0," *Natl. Inst. Stand. Technol.* **NISTIR 637** (1999).

[193] N. D. Mermin, "The topological theory of defects in ordered media," *Rev. Mod. Phys.* **51**, 591 (1979).

[194] K. Nakamura, T., Kawamura, "The Frustration of Antiferromagnetic Bonds in the Screw Dislocation," *Phys. Status Solidi B* **122**, 141 (1984).

[195] I. E. Dzyaloshkinskii, "Domains and dislocations in antiferromagnets," *JETP Lett.* **25**, 98 (1977).

[196] B. A. Ivanov and V. E. Kireev, "Spin disclination in a layered antiferromagnet with a screw dislocation," *J. Exp. Theor. Phys. Lett.* **73**, 188 (2001).

[197] B. A. Ivanov, "Mesoscopic antiferromagnets: statics, dynamics, and quantum tunneling," *Low Temp. Phys.* **31**, 635 (2005).

[198] I. Sugiyama, N. Shibata, Z. Wang, S. Kobayashi, T. Yamamoto, and Y. Ikuhara, "Ferromagnetic dislocations in antiferromagnetic NiO," *Nat. Nanotechnol.* **8**, 266 (2013).

[199] Y. Ran, Y. Zhang, and A. Vishwanath, "One-dimensional topologically protected modes in topological insulators with lattice dislocations," *Nat. Phys.* **5**, 298 (2009).

[200] E. C. Aifantis, "The physics of plastic deformation," *Int. J. Plast.* **3**, 211 (1987).

[201] M. Kuzmina, M. Herbig, D. Ponge, S. Sandlobes, and D. Raabe, "Linear complexions: Confined chemical and structural states at dislocations," *Science* **349**, 1080 (2015).

[202] W. Callister and D. Rethwisch, *Materials science and engineering: an introduction*, in *Mater. Sci. Eng.* **94** (2007).

[203] P. J. McNally, "Techniques: 3D imaging of crystal defects," *Nature* **496**, 37 (2013).

[204] G. von Heijne, "Membrane-protein topology," *Nat. Rev. Mol. Cell Biol.* **7**, 909 (2006).

[205] J. Shea, J. N. Onuchic, and C. L. Brooks III, "Exploring the origins of topological frustration : Design of a minimally frustrated model of fragment B of protein A," *Proc. Natl. Acad. Sci.* **96**, 12512 (1999).

[206] G. Gafvelin and G. Von Hetjne, "Topological 'Frustration' in Multispanning E. coli Inner Membrane Proteins," *Cell* **77**, 401 (1994).

[207] J. D. Bryngelson, J. N. Onuchic, N. D. Socci, and P. G. Wolynes, "Funnels, pathways, and the energy landscape of protein folding: a synthesis," Condensed Matter; Chemical Physics, *Proteins* **21**, 167 (1995).

[208] D. Baker, "A surprising simplicity to protein folding.," *Nature* **405**, 39 (2000).

[209] W. L. Wang, O. V. Yazyev, S. Meng, and E. Kaxiras, "Topological Frustration in Graphene Nanoflakes: Magnetic Order and Spin Logic Devices," *Phys. Rev. Lett.* **102**, 157201 (2009).

[210] W. L. Wang, S. Meng, and E. Kaxiras, "Graphene NanoFlakes with Large Spin," *Nano Lett.* **8**, 241 (2008).

[211] T. Araki, M. Buscaglia, T. Bellini, and H. Tanaka, "Memory and topological frustration in nematic liquid crystals confined in porous materials," *Nat. Mater.* **10**, 303 (2011).

[212] M. Kleiber, M. Bode, R. Ravlić, and R. Wiesendanger, "Topology-Induced Spin Frustrations at the Cr(001) Surface Studied by Spin-Polarized Scanning Tunneling Spectroscopy," *Phys. Rev. Lett.* **85**, 4606 (2000).

[213] R. Ravlić, M. Bode, A. Kubetzka, and R. Wiesendanger, "Correlation of dislocation and domain structure of Cr(001) investigated by spin-polarized scanning tunneling microscopy," *Phys. Rev. B* **67**, 174411 (2003).

[214] G. Y. Chin, "Slip and Twinning Systems in Ceramic Crystals," in *Deform. Ceram. Mater.* (Springer US, Boston, MA, 1975).

[215] E. Fawcett, H. L. Alberts, V. Y. Galkin, D. R. Noakes, and J. V. Yakhmi, "Spin-density-wave antiferromagnetism in chromium alloys," *Rev. Mod. Phys.* **66**, 25 (1994).

[216] D. P. Osterman, K. Mohanty, and J. D. Axe, "Observation of the antiferroelectric order parameter in surface layers of $SrTiO_3$," *J. Phys. C Solid State Phys.* **21**, 2635 (1988).

[217] G. Grüner, "The dynamics of charge-density waves," *Rev. Mod. Phys.* **60**, 1129 (1988).

[218] E. Berg, E. Fradkin, S. A. Kivelson, and J. M. Tranquada, "Striped superconductors: how spin, charge and superconducting orders intertwine in the cuprates," *New J. Phys.* **11**, 115004 (2009).

[219] P. Henelius, T. Lin, M. Enjalran, Z. Hao, J. G. Rau, J. Altosaar, F. Flicker, T. Yavors'kii, and M. J. P. Gingras, "Refrustration and competing orders in the prototypical $Dy_2Ti_2O_7$ spin ice material," *Phys. Rev. B* **93**, 024402 (2016).

[220] K. A. Ross, T. Proffen, H. A. Dabkowska, J. A. Quilliam, L. R. Yaraskavitch, J. B. Kycia, and B. D. Gaulin, "Lightly stuffed pyrochlore structure of single-crystalline $Yb_2Ti_2O_7$ grown by the optical floating zone technique," *Phys. Rev. B* **86**, 174424 (2012).

[221] K. Baroudi, B. D. Gaulin, S. H. Lapidus, J. Gaudet, and R. J. Cava, "Symmetry and light stuffing of $Ho_2Ti_2O_7$, $Er_2Ti_2O_7$, and $Yb_2Ti_2O_7$ characterized by synchrotron x-ray diffraction," *Phys. Rev. B* **92**, 024110 (2015).

[222] H. M. Revell, L. R. Yaraskavitch, J. D. Mason, K. A. Ross, H. M. L. Noad, H. A. Dabkowska, B. D. Gaulin, P. Henelius, and J. B. Kycia, "Evidence of impurity and boundary effects on magnetic monopole dynamics in spin ice," *Nat. Phys.* **9**, 34 (2012).

[223] G. Sala, M. J. Gutmann, D. Prabhakaran, D. Pomaranski, C. Mitchelitis, J. B. Kycia, D. G. Porter, C. Castelnovo, and J. P. Goff, "Vacancy defects and monopole dynamics in oxygen-deficient pyrochlores," *Nat. Mater.* **13**, 488 (2014).