

Probing antiferromagnetic hysteresis on programmable quantum annealers

Elijah Pelofske^{1,2,3,*}, Pratik Sathe^{2,4}, Cristiano Nisoli^{2,3,4,5} and Frank Barrows^{2,3,5}

¹Information Systems & Modeling, *Los Alamos National Laboratory*, Los Alamos, New Mexico 87545, USA

²Theoretical Division, Quantum & Condensed Matter Physics, *Los Alamos National Laboratory*, Los Alamos, New Mexico 87545, USA

³Center for Quantum Computing, *Los Alamos National Laboratory*, Los Alamos, New Mexico 87545, USA

⁴Information Science and Technology Institute, *Los Alamos National Laboratory*, Los Alamos, New Mexico 87545, USA

⁵Center for Nonlinear Studies, *Los Alamos National Laboratory*, Los Alamos, New Mexico 87545, USA



(Received 3 December 2025; accepted 30 September 2026; published 21 October 2026)

Using programmable analog quantum annealing processors, we implement a sampling-based magnetic hysteresis protocol to probe the counterintuitive notion of magnetic memory in antiferromagnetic models. A key component of this protocol responsible for the hysteresis is a transverse field, which enables state transitions, while the longitudinal magnetic field sweep is done via a longitudinal control field. We observe intriguing nonmonotonic magnetization, alongside full saturation and reversal of the hysteresis curve, Barkhausen-like noise, as well as emergent magnetic domains mediated by quantum fluctuations that give rise to the magnetic memory effect in antiferromagnets.

DOI: [10.1103/mcz4-xh6z](https://doi.org/10.1103/mcz4-xh6z)

I. INTRODUCTION

Quantum annealing is a type of analog quantum computation that is based on the adiabatic theorem of quantum mechanics [1–4]. While originally designed to solve optimization problems, quantum annealers are also emerging as programmable laboratories for physical phenomena which are otherwise inaccessible using either bulk material observables in magnetic laboratories, or classical simulation techniques [5–15].

In recent work we expanded the phenomenological range of these experiments by introducing a protocol to emulate magnetic hysteresis experiments using D-Wave’s quantum annealers, which are conceptually identical to measurements performed in a laboratory [16]. In Refs. [16,17] we applied this method to ferromagnetic and $\pm J$ disordered Ising models. In this study we apply our method to simulations of a counterintuitive type: hysteresis in antiferromagnetically interacting spins. We perform these experiments on four different D-Wave quantum annealing quantum processing units (QPUs). We find several interesting phenomena, including nonmonotonicities similar to those observed in low-dimensional ferromagnetic models [17]. Where meaningful, we extract both spin structure factor and the antiferromagnetic (Néel) order parameter from the two-dimensional (2D) antiferromagnetic models to probe the magnetic order throughout the hysteresis cycles.

Historically, magnetic hysteresis has most commonly been observed and studied in ferromagnetic compounds, where it arises from the breakdown of ergodicity associated with

the spontaneous symmetry breaking around a specific magnetization. While a sufficiently strong field can polarize an antiferromagnet, one would not easily expect hysteretic magnetization, because the local field acts to destroy the net moment.

Antiferromagnetic hysteresis is a rare but known phenomenon explained with diverse mechanisms in various systems [18–23], though typically not in *purely* antiferromagnetic models, where it remains poorly understood. Indeed, even the textbook explanation for standard, ferromagnetic hysteresis in terms of sticky metastabilities alone proves deceptively simple: glossing over, in the classical case, the effects of kinetics in dynamical hysteresis [24,25]; and neglecting, in quantum systems, the role of quantum fluctuations in driving state transitions via tunneling. This motivates a study of hysteresis precisely where such metastability is absent [25]: the pure antiferromagnet.

In this work we explore an exotic form of magnetic hysteresis in various systems of antiferromagnetically interacting qubits exposed to quantum fluctuations induced by transverse fields. Of the three varieties of models we consider, two exhibit frustration [26]. There exist many magnetic compounds that can be effectively modeled as antiferromagnetic quantum Hamiltonians with various quasi-one-dimensional (1D) or triangular (2D) lattices, such as $\text{BaCo}_2\text{V}_2\text{O}_8$ [27], $\text{SrCo}_2\text{V}_2\text{O}_8$ [28], and TmMgGaO_4 [29]. In this study, we focus on simple transverse-field Ising model open quantum systems that are 1D and 2D as canonical baseline geometries on which to probe antiferromagnetic hysteresis. We also experimentally show results from D-Wave processor lattices whose connectivity does not correspond to simple low-dimensional quantum Hamiltonian models. However, this protocol can be applied to many other lattices and other quantum Hamiltonians besides what we report here, in particular for D-Wave quantum processors any Ising model that can be expressed as a subgraph of the hardware graph. Because magnetic hysteresis is a fundamental emergent phenomenon of

*Contact author: epelofske@lanl.gov




FIG. 1. Schematic of the magnetic hysteresis protocol. (a) The field ramps over time, as schematics (not actual energy scales or timescales) where we begin recording magnetization at the maximum-polarization longitudinal field denoted by the dashed purple vertical line; (b) the conceptual tracing of the average magnetization response to the applied longitudinal field H_z . (c), (d) The programmed D-Wave hardware waveforms for the *final* simulation time of 11.2 μs , where s is the “anneal-schedule” control parameter, and the h -gain field is H_z . Adapted from Ref. [16].

interacting magnetic spins, the possibility of analog quantum experiments of antiferromagnetic hysteresis is a notable computational capability of contemporary, noisy, analog quantum annealing processors.

II. METHODS

A. The hysteresis protocol

We perform analog simulations on programmable quantum annealers manufactured by D-Wave Systems [30,31] (see Fig. 1). These devices use superconducting flux qubits [32–36] connected in a graph comprising repeating subgraph units, and operate at approximately 15 millikelvin (15 mK). D-Wave hardware realizes time-dependent transverse field Ising models of the form

$$\mathcal{H} = -\frac{A(s)}{2} \sum_i \hat{\sigma}_x^{(i)} + \frac{B(s)}{2} \left[g(t) \sum_i h_i \hat{\sigma}_z^{(i)} + \sum_{i \neq j} J_{ij} \hat{\sigma}_z^{(i)} \hat{\sigma}_z^{(j)} \right], \quad (1)$$

where $s \in [0, 1]$ is a programmable time-dependent annealing parameter, and $A(s)$ and $B(s)$ are hardware-dependent energy scales (see Appendix A). The dimensionless coupler interaction strengths $\{J_{ij}\}$ and longitudinal fields $\{h_i\}$ can be programmed by the user. The σ_x Pauli term is the noncommuting driving Hamiltonian. Crucially, the function $g(t)$, which multiplies the longitudinal field, can be controlled independently of s , and is thus central to our methods, as described in Ref. [16]. In all our experiments, we set each h_i to the maximum allowed normalized-hardware programmable value.

The ability to vary g in time, central to our experiments, is currently available only in the total-simulation-time regime of on the order of microseconds. While quantum annealers behave coherently in the fast annealing regime (with anneal times of the order of tens of nanoseconds) [37,38], in the slow anneal regime, the devices are described more accurately by an open quantum system [39]. In this regime, coupling to the environment leading to thermalization [40–43], analog control errors and cross talk [44–47], time-dependent noise characteristics [48], and spin bath polarization [49] can all cause deviations from a closed quantum evolution. For all measurements we use an 11.2 microsecond total annealing time. We implement “forward anneals” (i.e., those where s is a

TABLE I. D-Wave QPU summary, where the maximum possible h -gain field that can be applied is in hardware-normalized units (relative to the specific Hamiltonian on each QPU).

D-wave QPU chip	Graph name	Qubits	Couplers	Maximum h -gain field strength
Advantage_system4.1	Pegasus P_{16}	5627	40279	± 1.75
Advantage_system6.4	Pegasus P_{16}	5612	40088	± 4
Advantage_system7.1	Pegasus P_{16}	5554	39238	± 3.5
Advantage2_prototype2.6	Zephyr $Z_{6,4}$	1248	10827	± 3

nondecreasing function of time, which starts and ends at 0 and 1 respectively) with a pause in the annealing parameter s during which we activate the longitudinal field. As in Ref. [16], we report average magnetization as a function of applied longitudinal field, with a reversed magnetization sign to compensate for the sign of the D-Wave QPU Hamiltonian Eq. (1).

This hysteresis protocol is a probabilistic, sampling based protocol, and it requires both the transverse field (denoted as Γ) on the D-Wave QPUs, which facilitates state transitions, and a sufficiently strong lattice coupling J , without which no hysteresis can be measured. Our protocol has several important tunable parameters. The physical energy scale ratio Γ/J at which the hysteresis cycle is swept has distinct behaviors at the two extremes. If it is too small, the state does not change (not even due to thermal fluctuations). On the other hand, if Γ/J is too large, then we are left with featureless unpolarized states with memory-less trajectories. Additionally, the longitudinal field must be sufficiently strong to force the system to become fully polarized against the antiferromagnetic bonds: therefore if the D-Wave hardware controls do not allow significantly strong longitudinal fields, only minor hysteresis loops are obtained. In principle one could always obtain saturation by using smaller coupling values, but that would amplify the role of thermal fluctuations and further destroy memory in an undesirable manner. Each QPU also has a different maximum longitudinal field that can be applied; the maximum field values of the “ h -gain” schedule, i.e., of $g(t)$ in Eq. (1) are provided in Table I. With the exception of the quasi-3D QPU hardware graph experiments, for all the reported experiments, we use default hardware settings, with no additional analog tuning; for instance, the flux bias offsets were left to default values, and no gauge averaging was applied.

B. Models and embeddings

We consider three different antiferromagnetic models: 1D with periodic boundary conditions, 2D with open boundary conditions, and a specific type of “pseudo-three-dimensional” (pseudo-3D) model that involves the full hardware graph. For all models, we set the coupler J values to 1, the maximum coupler strength that can be specified on the hardware.

The 2D square-grid antiferromagnetic model case corresponds to entirely unfrustrated magnetic systems because they have open boundary conditions. The pseudo-3D model is magnetically frustrated; it is not bipartite. The 1D antiferromagnetic systems we study have an odd number of spins and are hence frustrated on one bond.

For embedding or mapping the 1D and 2D models onto the hardware graph, we used the Glasgow solver [50], a graph isomorphism finder, which is a part of the `minorminer` package

[51]. The maximum system sizes for the 1D model we were able to embed are: 4905 qubits on Advantage_system4.1; 4885 qubits on Advantage_system6.4, 4791 qubits on Advantage_system7.1; and 1131 qubits on Advantage2_prototype2.6. For the 2D model, the maximum sizes of the models we realize on the QPU are 32×32 lattice on Advantage_system4.1 and Advantage_system6.4, 33×33 on Advantage_system7.1, and 26×26 on Advantage2_prototype2.6. See Appendix H for visual renderings of the embeddings.

In the hardware-defined antiferromagnetic models, we set every programmable coupler on the QPU to $J = 1$, which is the maximum coupler strength that can be specified on the hardware (in normalized programmable units) and whose energy scale is given in Appendix A. The number of qubits and couplers for the various QPUs are listed in Table I. The QPUs have either Pegasus [52,53] or Zephyr [54] connectivity.

C. Calibration refinement by statistical balancing

Nonidealities in the quantum annealers, even small fluctuations in the magnetic environment, can result in nonuniform response even in models that are uniform. These effects can be mitigated by leveraging the symmetries of the model being studied via an iterative statistical-balancing calibration refinement technique [55], which has been implemented in a broad range of experiments [10,12,15,37,38,55–59]. In some of our “quasi-3D” hardware-defined experiments, we implement a flux bias offset (FBO) calibration refinement procedure. In particular, before implementing any hysteresis experiment, we implement trial experiments with zero longitudinal field. Due to $\mathbb{Z}_2$ symmetry, the average magnetization of each qubit should be 0. Due to experimental nonidealities, we instead find a spread in the values of qubit magnetizations. By iteratively modifying the FBOs, the spread of qubit magnetizations can be reduced, achieving more uniformity (see Appendix B for more details). The FBO values obtained at the end of the calibration refinement procedure are then used for the hysteresis experiments to improve accuracy of the results. While more advanced calibration refinement techniques could be utilized, such as modifying coupler strengths to make coupler frustration uniform, or refining individual h_i values in each “time slice” of the hysteresis protocol, such a procedure would use significantly more QPU time particularly for models of the size we study here. Hence, we restrict to FBO refinement in this work. For completeness we report both calibrated and uncalibrated data for the quasi-3D models.




FIG. 2. Magnetic hysteresis on 1D odd antiferromagnetic rings. Average magnetization as a function of the applied longitudinal field H . The overlaid black arrows denote the time progression of the experiment as well as the direction of the longitudinal field sweeps. The different lines correspond to D-Wave experiments performed at different anneal-schedule pauses, defined by s , which corresponds to a particular physical field ratio of Γ/J .

D. Observables

From the experimental data, we extract the average longitudinal magnetization

$$M^z = \frac{1}{N} \sum_i \langle \sigma_i^z \rangle, \quad (2)$$

where $\langle \dots \rangle$ denotes an average over 2000 samples (8000 samples in the case of the quasi-3D models) measured on the hardware sampler. The experimental average over a relatively large number of samples mitigates the possibility of apparent ordering due to finite sampling effects, which may occur if the number of samples is especially small.

It is natural to also consider the antiferromagnetic order parameter, which is defined as

$$M_s = \frac{1}{N} \sum_i (-1)^i \langle \sigma_i^z \rangle. \quad (3)$$

This quantity is well defined only for a bipartite graph, and hence can be extracted only for the 2D models and not the odd-length 1D chains or the quasi-3D models. Interestingly, unlike magnetization, the *staggered* antiferromagnetic order parameter is inaccessible to experiments in a magnetic laboratory without a scattering experiment such as

neutron scattering, but is computationally easily accessible to experiments on quantum annealers because they provide information on the elementary degrees of freedom.

Our access to the individual degrees of freedom allows for the direct extraction of other interesting observables, including the real space snapshots of spin configurations, which we will show in Sec. III. This also allows us to extract the magnetic structure factor, which gives clear visual evidence for different types of magnetic ordering during hysteresis cycles. While in experiments it is obtained by neutron scattering, we can reconstruct it from individual spins, as

$$S(\mathbf{q}) = \left\langle \sum_{i,j} e^{i\mathbf{q} \cdot (\mathbf{r}_i - \mathbf{r}_j)} \sigma_z^i \sigma_z^j \right\rangle, \quad (4)$$

where $\mathbf{r}$ denotes spin lattice locations. For the 2D case on a square lattice, we discretize the momentum vector $\mathbf{q}$ on a 200×200 uniformly spaced grid spanning $(-2\pi, 2\pi)$ along both dimensions.

The pseudo-3D models lack a sensible definition for a Brillouin zone, and hence we compute the structure factor only for the 1D and 2D models.

For the 1D models (which always have odd lengths in our study), we examine domain wall counts which ranges from 1




FIG. 3. Selected representative example 1D antiferromagnet spin configurations measured during the hysteresis cycle run on Advantage2_prototype2.6 (at $s = 0.7$), comprising 1131 spins. Each spin configuration is from a different applied longitudinal field value in the nonsaturated region of the hysteresis cycle, both near $H \approx -2$. The spin configurations are shown as a circular annular wedge plot where black slices denote spin-down qubit measurements and cyan slices denote spin up qubit measurements. The net magnetization for each sample is shown in the center of each plot. These selected samples show different cases of fragmented antiferromagnetic ordering, which is to be expected for these experiments.

(in the ground state) to L for the fully polarized state (where L is the number of sites in the 1D model).

III. RESULTS

We report antiferromagnetic hysteresis for the 1D, 2D, and quasi-3D models in Secs. III A, III B, and III D respectively. All reported results for the 1D and 2D (quasi-3D) models are computed from 2000 (8000) measurements per data point, for each change in experimental configuration. All reported net magnetization hysteresis results are statistical averages. The hysteresis cycle plots are presented as a function of H , where H field units are in hardware-specific programmable units.




FIG. 4. Domain wall density on a 1D antiferromagnetic ring during the hysteresis cycles as a function of applied longitudinal field H . Left-hand plot shows the total, average, domain wall density, regardless of the sign of the domain wall. Right-hand plot shows the spin-down ($\downarrow\downarrow$) domain wall density. Each closed curve corresponds to a different Γ/J value on the D-Wave processor. Note that because these rings have an odd number of spins, we are always guaranteed to have at least one pinned domain wall somewhere on the ring. Color encodes the Γ/J energy scale ratio, which corresponds to an “anneal”-pause value s . Both at very high Γ/J and small Γ/J the antiferromagnetic domain wall proportion drops to very close to zero, which is due to state transitions being energetically favorable in one case and unfavorable in the other case. However, the lack of domain walls means that for most of the plotted color spectrum, the data lines are nearly identical, and therefore visually only a certain range of Γ/J show differences.

A. 1D antiferromagnet experiments

Figure 2 presents magnetic hysteresis experiments for 1D antiferromagnetic models with periodic boundary conditions, performed at different values of Γ/J , which we control by varying the value of s . It shows interesting nonmonotonic features. In particular, a “pinch-point” occurs at $s = 0.6$ on three of the QPUs, except Advantage2_prototype2.6. When Γ/J is sufficiently small, we achieve full polarization, unlike for larger values of Γ .

Figure 3 shows examples of real-space spin configurations sampled during one of the hysteresis cycles run on the D-Wave QPU Advantage2_prototype2.6, specifically in the nonsaturated regions of the cycle. These spin configurations show very small antiferromagnetic domains at small M_z ; the distribution of alternating sign spins is not uniform, which could be in part due to small J coupling energy scales differences on the analog hardware. One sees formation of magnetic domains for larger M_z , as expected.

Figure 4 reports 1D domain wall density in the 1D hysteresis cycles, including the proportion of $\downarrow\downarrow$ domain walls (note that, like magnetization, the sign of this domain wall has been reversed in order to compensate for the sign of the D-Wave processor Hamiltonian [16]). The domain wall density on the other three QPUs is qualitatively similar, which can be found in Appendix E.

1. Odd rings, domain-wall sectors, and their use as a defect transport probe

To provide more insight into these experiments, consider the transverse-field Ising antiferromagnet on a ring,

$$H = J \sum_{i=1}^N \sigma_i^z \sigma_{i+1}^z - h \sum_{i=1}^N \sigma_i^z - \Gamma \sum_{i=1}^N \sigma_i^x, \quad (5)$$

with $\sigma_{N+1}^z \equiv \sigma_1^z$ and $J > 0$. Because N is odd, there are an odd number of domain walls for any computational basis (σ^z basis) configuration. When $|h|, \Gamma \ll J$, the ground-state subspace, comprising configurations with exactly one domain wall ($D = 1$), is N -fold degenerate.

Local spin flips generated by the transverse part of the Hamiltonian, proportional to Γ , toggle the two bonds adjacent to the flipped site, so the total number of domain walls changes by 0 or ± 2 , thus conserving the parity of domain wall count. Thus, when $h = 0$, domain wall parity is a conserved quantity. Additionally, the domain walls come in two flavors, corresponding to both the spins pointing up, or down. A longitudinal field $h \neq 0$ breaks the global $\mathbb{Z}_2$ symmetry and splits these two local configurations in energy. These are not distinct dynamical sectors once $\Gamma > 0$, because a single spin flip at the wall converts a spin-up domain wall to a spin-down domain wall (and vice versa) while shifting the wall by one bond. Thus, odd rings naturally realize a single-wall system in which Γ drives coherent domain-wall motion and h biases its drift.

This built-in, mobile defect makes odd rings a probe of quantum-assisted defect dynamics and antiferromagnetic hysteresis. In contrast, an even- N ring possesses a wall-free Néel lowest energy configuration; reversal at small $|h|$ must first nucleate a kink–antikink pair (energy cost $\approx 4J$ at $\Gamma = 0$), which strongly suppresses response and complicates interpretation. Because even rings require pair creation, their low-field loops are comparatively featureless and are not informative about defect-transport dynamics, we therefore concentrate on odd rings for quantitative analysis. For this reason, and because our experiments target quantum assisted transport and pinning phenomena, we restrict to odd rings in all the 1D experiments in this study.

Weak static inhomogeneity (bond or field offsets) pins the wall at preferred bonds. At a pinned location the two local sign states form a driven two-level subsystem; the transverse field Γ mixes them, while h sets their energy difference. As detailed below, this competition produces an avoided crossing and underlies the stick-slip dynamics and steplike magnetization changes characteristic of the antiferromagnetic hysteresis we observe.

In summary, the 1D antiferromagnet model experiments show nonmonotonic magnetization features and interesting switching that occurs at high longitudinal field. These features have been observed before in prior quantum annealer experiments on different types of models [16,17]. In these models however, there is more “corrugation” compared to ferromagnet models [17]. Examining domain wall statistics and individual measurements reveals fragmented antiferromagnetic ordering in the approximately depolarized region of the hysteresis cycles.

2. Domain-wall theory: Effective Hamiltonian, dispersion, and many-wall processes

In the single-wall regime we label basis states $|j, w\rangle$, where $j \in \{1, \dots, N\}$ denotes the location of the domain wall and $w = \pm 1$ denotes its flavor (both up or both down). For these states, the exchange energy contribution is given by

$$E_J = -J(N - 2), \quad (6)$$

because one bond is ferromagnetic ($+J$) and the remaining $N - 1$ bonds are antiferromagnetic ($-J$). The longitudinal field contributes a term that depends only on the wall sign: for odd N the total magnetization in a one-wall state is

$$M = \sum_i \sigma_i^z = w, \text{ so}$$

$$H_h \rightarrow -hw \quad \text{on} \quad |j, w\rangle, \quad (7)$$

where the longitudinal-field term is denoted by $H_h = -h \sum_i \sigma_i^z$. A single σ^x acting on either spin adjacent to the wall reverses that spin, toggles the wall sign, and shifts the wall by one bond. To leading order in Γ/J , degenerate perturbation theory yields an effective Hamiltonian given by

$$H_{\text{eff}} = E_J + \sum_{w=\pm 1} \sum_{j=1}^N [-hw |j, w\rangle \langle j, w| - \Gamma(|j+1, -w\rangle \times \langle j, w| + |j-1, -w\rangle \langle j, w|)], \quad (8)$$

with periodic indexing $j \equiv j+N$. Equation (8) is a tight-binding model on a two-sublattice (sign) chain.

For a homogeneous ring we diagonalize with Bloch states $|k, w\rangle = \frac{1}{\sqrt{N}} \sum_j e^{ikj} |j, w\rangle$, $k = 2\pi m/N$. In the $\{|k, +\rangle, |k, -\rangle\}$ basis the Bloch Hamiltonian reads

$$\mathcal{H}(k) = -h\sigma^z - 2\Gamma \cos(k)\sigma^x, \quad (9)$$

with eigenenergies

$$E_{\pm}(k) = E_J \pm \sqrt{h^2 + 4\Gamma^2 \cos^2(k)}. \quad (10)$$

Thus, the transverse field endows the wall with a hopping amplitude $\sim \Gamma$ and, at $h \approx 0$, a single-domain-wall spectrum spanning an energy width 4Γ ; the time-dependent longitudinal field can populate excited one-wall states. A finite longitudinal field splits the two bands by at least $2|h|$ (at $k = \pi$), producing a drift bias for wall motion. The corresponding group velocity is

$$v_{\pm}(k) = \pm \frac{2\Gamma^2 \sin(k)}{\sqrt{h^2 + 4\Gamma^2 \cos^2(k)}}. \quad (11)$$

While odd N fixes the wall-number parity, additional kink–antikink pairs can be nucleated at finite drive. At $\Gamma = 0$ and small $|h|$ a single local spin flip in a Néel region creates two additional domain walls and costs

$$\Delta E_{\text{pair}} \approx 4J - 2h\delta M, \quad (12)$$

with $\delta M = \pm 1$ in our normalization ($\sum_i \sigma_i^z = M$). Quantum mechanically, pair creation arises from higher-order virtual processes in Γ/J , while annihilation occurs when a wall meets an antiwall. On our platform thermal activation is negligible; the dynamics is therefore dominated by wall hopping, field-biased drift, quantum-assisted pair nucleation, and annihilation.

The interplay of hopping ($\sim \Gamma$) and field bias (h), pinning, and rare nucleation and annihilation events produces the antiferromagnetic hysteresis observed in odd rings. In the single-wall-dominated regime, magnetization evolves via sequences of hops punctuated by depinning events at local avoided crossings, generating discrete, Barkhausen-like steps. As the drive amplitude increases, occasional nucleation events seed additional walls which drift and annihilate; this “dilute domain-wall gas” smooths the loop while preserving a strong odd-ring response near low fields. A simple estimate gives the pair-creation energy $\Delta E_{\text{pair}} \approx 4J - 2h\delta M$, so pair nucleation




FIG. 5. Magnetic hysteresis on 2D antiferromagnetic square lattices, on four different D-Wave quantum annealers. Average magnetization (y axis) as a function of the applied longitudinal field (x axis). The overlaid black arrows on the average magnetization lines sampled from the QPUs simultaneously denote the time progression of the analog simulations, as well as the direction of the longitudinal field sweeps. The lower-right subplot results do not get close to saturation for any Γ/J energy scale ratio causing the magnetization lines to overlap. The reason for this is because the maximum longitudinal field hardware energy scales are not as large as for the other processors.

becomes likely once $|h| \sim 2J$, beyond this “nucleation threshold” the loop evolves in a dilute domain-wall gas rather than by single-wall motion. In the data the rise of total domain wall density and magnetization reversal set in around this scale.

For odd- N antiferromagnetic rings the domain-wall number D is constrained to be odd, in the regime of small $|h|$ with $\Gamma \ll J$ the lowest-energy states have a single domain wall. In Fig. 4, in the same low-field window the sign-resolved wall proportion exhibits a strong, hysteretic bias between $\uparrow\uparrow$ and $\downarrow\downarrow$ walls, with forward–reverse asymmetry. This pattern—weak change in total domain wall proportion but strong, history-dependent redistribution between the two wall signs—is exactly what the effective single-wall Hamiltonian predicts: h splits the two wall-sign states $w = \pm 1$ while Γ mixes them, yielding a biased two-level population that traces a hysteresis loop. In equilibrium this bias would reduce to a sigmoidal form $\frac{1}{2}[1 + \tanh(\beta_{\text{eff}}h)]$.

As $|h|$ grows the total domain wall proportion becomes large and the loop area is dominated by many-wall dynamics. The conspicuous nonmonotonicity near $|h| \approx 1$ at $s \approx 0.6$ lies in this regime. Such large nonmonotonicities can be understood as a result of synchronized depinning at reproducible nucleation fields set by hardware pinning, i.e., pinning sites

and hopping barriers introduced by hardware disorder and inhomogeneities in device parameters. The nonmonotonicities and resulting Barkhausen steps survive ensemble averaging because those nucleation fields recur shot to shot.

B. 2D antiferromagnet experiments

Figure 5 presents 2D antiferromagnetic hysteresis simulations. Again we observe full saturation for some values of Γ/J , on most QPUs. On the Advantage2_prototype2.6 processor, the maximum longitudinal field that can be applied is comparatively weaker, resulting in the net magnetization being lower than ± 0.5 for any Γ/J value. In some regimes, such as $s = 0.7$ and $s = 0.75$ where the system is fully polarized, we observe Barkhausen-like noise (e.g., crackling) in these experiments.

Here we can also extract the spin structure factor in order to examine the types of magnetic ordering that occur during the hysteresis cycles. Figure 6 shows averaged magnetic structure factors¹ from various points along the antiferromagnetic hys-

¹Also referred to as Static Structure Factor (SSF).




FIG. 6. Magnetic structure factors $|S(q)|$ heatmaps (bottom row) within the 2D antiferromagnetic hysteresis cycles (averaged over 100 independent samples), run on Advantage_system6.4 at $s = 0.7$. The Ising model being simulated is a 32×32 grid of (antiferromagnetically coupled) spins, open boundary conditions. The top row shows the single hysteresis cycle in terms of M_z average magnetization, from which we extracted $|S(q)|$ heatmaps for four points along the first longitudinal field sweep, denoted by cyan asterisks. The dashed white line on the structure factor heat maps outlines the first Brillouin zone. The heat maps are log-scaled and scales are specific to each heatmap subplot. The order of the $|S(q)|$ heatmaps follows the same sweep direction denoted by the top hysteresis cycle plot (the static structure factor from the fully saturated regime is shown in the right hand subplot). Note the peaks at the corner of the Brillouin zone denoting large antiferromagnetic clusters, and coexisting with a growing peak at the center denoting magnetization.

teresis cycle. We see clear boundary effects (due to the open boundary conditions) that occur closer to the demagnetized regions, seen as lines along $x = y = 0$ axes. In the saturated regimes, we see either bright peaks along $q_x = \pi$ and $q_y = \pi$ when Γ/J is high (meaning, the magnetization is not fully saturated), or closer to ferromagnetic-like ordering (only due to the system being polarized) when the ratio Γ/J is smaller and also the longitudinal field is strong enough. The clearest signal we see is strong antiferromagnetic ordering, with peaks at the four corners of the first Brillouin zone, as the systems become more demagnetized. Across the different D-Wave QPUs and physical device settings, there are slight differences in the magnetic ordering captured by the SSF plots, but qualitatively they are all similar to each other with the primary difference being to what extent full saturation can be achieved; additional SSF plots under other settings are shown in Appendix C.

Figure 7 shows examples of real space single-spin configurations measured during one of the 2D antiferromagnet hysteresis cycles. These samples show fragmentation into domain walls of the expected checkerboard ordering in the ground state of the 2D square antiferromagnet, and moreover we see clear boundary effects (due to the open boundary conditions) as the system becomes more magnetized. Qualitatively, the measured spin configurations during the hysteresis cycles are similar across the different devices and different physical energy scales, Appendix G reports similar measurements on a different D-Wave QPU. To better illustrate the antiferromagnetic domains in Fig. 7 we plot also on the bipartite lattice the domain height function

$$\tau_{i_x i_y} = (-1)^{i_x + i_y} \sigma_{i_x i_y}^z, \quad (13)$$

which is uniform and valued ± 1 in each domain. Domain walls are then the contours separating $\tau = +1$ and $\tau = -1$.

Clearly, it is related to the antiferromagnetic order parameter via $M_s = \sum_{i_x i_y} \tau_{i_x i_y} / N$.

In our finite 2D samples the reversal appears to proceed through a competition between nucleation of favorably oriented antiferromagnetic droplets (e.g., antiferromagnetic domains, seen as contiguous cyan or black regions in the lower row of Fig. 7) and propagation and roughening of preexisting interfaces, associated with antiferromagnetic symmetry breaking, as shown by the staggered boundaries (cyan-black pixels) between domains in the lower row of Fig. 7. The antiferromagnetic domains, and the boundaries between those domains, can be seen visually by the checkerboard ordering in the spin configurations in the top row of Fig. 7, but are better visually highlighted by the staggered height function in the bottom row of Fig. 7. Open boundaries play a decisive role: edges and, in particular, corners reduce the interfacial cost, so the first irreversible events predominantly originate at the boundary. Once a supercritical droplet appears, growth continues mainly by interfacial translation (with occasional coalescence when neighboring droplets meet), producing the steplike changes in magnetization that are characteristic of pinned, driven interfaces. At larger field magnitude the rate of new-droplet formation increases and the loop correspondingly smooths as propagation competes with fresh nucleation. In the raw-spin plots of Fig. 7 (top), the open-boundary effect appears as extended runs of similarly aligned spins along the outer edges, where spins have fewer antiferromagnetic neighbors than in the bulk. In the height-function plots of Fig. 7, changes across the sampled field values highlight antiferromagnetic domains (cyan and black pixel regions) and in the shapes of their contours are consistent with propagation and roughening of the corresponding domain walls.




FIG. 7. Selected real-space spin configurations sampled on the Advantage_system6.4 QPU, illustrated as pixel plots where black pixels denote spin-down qubit measurements and cyan pixels denote spin-up qubit measurements (top row), from the middle of the magnetization reversal at various longitudinal field strengths of the first magnetic hysteresis sweep of a 32×32 2D square grid antiferromagnet, at $s = 0.7$. Notice that the checkerboard antiferromagnetic ground-state emerges in sections of the lattices. Also notice that the boundary effects from the open boundary conditions are noticeable in many of the measured spin configurations. Bottom row: staggered domain height function τ from Eq. (13), corresponding to the same spin configurations on top row, making the antiferromagnetic domains apparent.

These processes leave clear fingerprints in the spin-structure factor. Near saturation the structure factor weight concentrates around $\mathbf{q}=0$, reflecting the large polarized regions. As the field is swept toward the demagnetized regime, intensity migrates to antiferromagnetic wavevectors at the corners of the Brillouin zone, consistent with the build-up of checkerboard order driven by local field, inside expanding droplets and along translating interfaces. As the field sweeps, the hysteresis emerges from the polarized domain walls.

Because of the finite window and open boundaries, we also observe streaks of intensity aligned with the axes of the Brillouin zone at $q_x = q_y = 0$, which track the prevalence of straight boundary-pinned segments. The streaks of intensity at $q_x = q_y = 0$ can be seen in all four structure factor panels of Fig. 6, but most clearly in the lower right panel which shows ordering when the model is entirely saturated. On the return sweep, the same features recur with a field offset set by the pinning landscape, yielding a reproducible hysteretic trajectory in both the net magnetization and the structure factor.

While both reciprocal space and real space data show formation of large antiferromagnetic domains and their coexistence with net magnetization, a relevant question is whether the data show order that can be properly interpreted as spontaneous antiferromagnetic symmetry breaking. To explore that, we move to the Néel order parameter [Eq. (3)] in Appendix F, which reveals that there is some latent bias in these 2D experiments that should be mitigated in future work using analog control statistics balancing methods.

In summary, these 2D experiments show interesting non-monotonic magnetization features, the models can be fully saturated given appropriate tuning of Γ/J and the longitudinal field is sufficiently strong. Moreover, the extracted static structure factors show ordering that is consistent with an antiferromagnet model undergoing reversal and saturation at high longitudinal field. By extracting individual spin configurations we can visually see antiferromagnetic domains, as well as

boundary effects due to the open boundary conditions of the model embedded on the hardware.

C. Hysteresis curve areas from 1D and 2D antiferromagnetic hysteresis cycles

Figure 8 plots the numerically integrated enclosed area of the hysteresis cycles as a function of s (equivalently, Γ/J), from the 1D and 2D antiferromagnet simulations. This shows that the region of the annealing parameter (s , equivalently Γ/J) where the hysteresis loop area is nonzero is smaller compared to the ferromagnetic models or disordered frustrated models studied in Ref. [16].

Across 1D rings and 2D grids the hysteresis loop area, as a function of s , is nonmonotonic in s : it turns on from near zero at small s , peaks in an intermediate window, and decreases again at larger s . This reflects a competition between pinning-limited reversal at low Γ/J and strong transverse-field mixing at high Γ/J , with device-to-device amplitude set primarily by the usable longitudinal field range, seen in Table I. The reduced longitudinal field produces minor loops and systematically smaller areas at the same s . In two dimensions, areas are further suppressed relative to one dimension at matched s , consistent with reduced magnetic saturation in the 2D case.

Statistical uncertainty quantification, measured by magnetization standard deviation, is reported in Appendix D.

D. Quasi-3D antiferromagnet experiments

Figures 9 and 10 show magnetic hysteresis simulations performed on 3D antiferromagnets. Specifically these use the entire hardware graph (every coupler being set to an antiferromagnetic coupling).

Figure 9 shows hysteresis curve data from Ising models that are *entirely antiferromagnetic*. These simulations are notable because while full magnetization is not reached on any of the antiferromagnetic lattices, the nonzero transverse




FIG. 8. Magnetic hysteresis cycle area (y axis) versus the annealing parameter s at which the hysteresis is performed, for the 2D square antiferromagnetic lattice with open boundaries, (left) and 1D antiferromagnetic rings, with periodic boundary conditions (right).

field does allow state transitions to take place and we recover a minor-loop hysteresis curve that resembles classical ferromagnetic hysteresis. Moreover, like in the prior antiferromagnetic models we observe Barkhausen-like steps in the magnetization cycles (Figs. 9 and 10), however, the apparent strength of these signals is weak because of the high shot noise and low-magnetization values. Being able to probe antiferromagnetic hysteresis using this technique is a notable

computational capability given the relatively high coordination number of these graphs.

The general trend that we observe here is that as the coordination number of the antiferromagnetic Ising models increases, it becomes harder to fully magnetize the systems in the analog QPU simulation. This is to be expected, because a higher coordination number means more energy is required,




FIG. 9. Magnetic hysteresis of 3D hardware-defined antiferromagnets, with no calibration refinement. Average magnetization (y axis) as a function of the applied longitudinal field (x axis). Because the absolute range of magnetization (y axis) is much smaller than the other Ising models, here we increase the shot count (i.e., the total number of samples obtained by a complete anneal-readout cycle) to 8000 for each data point. This increase in the number of samples reduces local jitter of the magnetization curves. No flux bias offset calibration was used for these simulations (e.g., before calibration).




FIG. 10. Magnetic hysteresis on hardware-defined antiferromagnets, with extensive whole-lattice FBO calibration refinement. Average magnetization (y axis) as a function of the applied longitudinal field (x axis). Each data point on the hysteresis curves are averaged single site magnetization from 8000 samples. In comparison to Fig. 9 where no statistical balancing calibration was used, here we see in general smoother magnetization curves, in particular on the Advantage_system7.1 processor.

in this case a stronger longitudinal field, in order to flip the spins against the strong J lattice coupling.

In summary, these quasi-3D antiferromagnet hysteresis experiments show that sufficiently high coordination number substantially suppresses the hysteresis effects. Despite this, some hysteresis can still be observed. Moreover, after applying observable bias reduction techniques the hysteresis effects remained effectively unchanged.

Suppressed hysteresis in 3D hardware lattice

Our 3D results exhibit only small, often near-minor loops. In highly connected antiferromagnetic graphs due to the large interfacial cost set by the coordination number and J , there is limited mobility when the transverse field Γ is either too small (frozen) or too large (paramagnetic).

On an average c -coordinated AF lattice/graph, any domain surface cuts a number of bonds proportional to c per unit area. Each cut bond contributes an energy scale $\sim J$, so the effective surface tension obeys

$$\lambda \propto cJ,$$

up to geometric factors set by the lattice or embedding. In three dimensions, the excess energy of a droplet, assumed

nearly spherical, of reversed Néel order with radius R is well approximated by a surface term that scales as

$$E_{\text{surf}}(R) \simeq 4\pi\lambda R^2.$$

Higher coordination therefore stiffens domain surfaces and penalizes both surface translation and the creation of new droplets.

A longitudinal field h favors spins aligned with its sign. Flipping a spin gains Zeeman energy $2|h|$ (in our $\sigma^z = \pm 1$ normalization), so a droplet containing a volume $\frac{4}{3}\pi R^3$ gains

$$E_{\text{field}}(R) \simeq -\frac{4}{3}\pi R^3(2|h|m_0), \quad (14)$$

where m_0 is the spin magnitude per site (for perfectly staggered order $m_0 = 1$). Setting the derivative of the sum of the surface and field terms to zero, $d(E_{\text{surf}} + E_{\text{field}})/dR = 0$, yields the critical radius

$$R_c = \frac{2\lambda}{2|h|m_0} = \frac{\lambda}{|h|m_0} \quad (15)$$

and an activation barrier

$$E_c = E_{\text{surf}}(R_c) + E_{\text{field}}(R_c) = \frac{4\pi}{3} \frac{\lambda^3}{|h|^2 m_0^2}. \quad (16)$$

Thus, in three dimensions the surface-area cost increases both R_c and E_c sharply, in addition the high hardware

coordination c raises λ and therefore pushes R_c and E_c up, suppressing bulk nucleation across the sweep range we employ.

On finite samples, droplets can nucleate as caps on faces (roughly halving both surface and volume terms) or at edges and corners (reducing cost further). This boundary-assisted nucleation produces small droplets and corresponding small hysteresis. Most field steps only rattle pinned facets without producing large, hysteretic changes in M_z .

The interfacial picture also clarifies how the control knobs set the loop morphology. First, the sweep amplitude (envelope) determines whether droplets can ever become supercritical: if the field never reaches values for which the critical radius R_c is microscopic, nucleation remains subcritical and the response collapses to minor, low-area cycles. Second, the sweep rate and dwell time control the opportunity for rare, multispin processes that assist nucleation and for marginal droplets to grow: faster ramps and shorter dwells reduce this opportunity and thus diminish loop area. Third, the transverse field fixes the matrix element for local surface moves and pair creation; for $\Gamma/J \ll 1$ surfaces are effectively frozen, while for very large Γ/J , strong quantum mixing erases memory of the sweep path so that hysteresis collapses. A pronounced, reproducible loop therefore occurs only in an intermediate window of Γ/J that is large enough to enable surface translation across pinning sites but not so large as to randomize domains during the protocol.

IV. DISCUSSION AND CONCLUSION

We have demonstrated for the first time large-scale simulations of antiferromagnetic hysteresis using quantum computers. This includes observation of unexpected non-monotonic magnetization curves as has been seen before in these types of analog QPU simulations on 1D ferromagnets [17], along with extraction of the spin structure factor, which shows clear transitions of magnetic ordering during the hysteresis cycles. The structure factor plots show a clear antiferromagnetic ordering transition from ferromagnetic like ordering when the systems are fully polarized, as well as boundary effects. We have also demonstrated that the use of statistics-balancing calibration on the analog quantum hardware can improve the quality of the hysteresis simulations, as demonstrated on full-hardware-graph antiferromagnets. One of the primary results that we demonstrate is that for low-coordination antiferromagnets (1D and 2D Ising models in particular), under an appropriately tuned ratio of transverse field (Γ) to lattice coupling (J), we can achieve full polarization of the antiferromagnets. Furthermore, the subsequent reversal resembles many of the key features of standard classical hysteresis in disordered, and even ferromagnetic, spin systems.

Our measurements establish robust dynamical hysteresis in programmable antiferromagnets across geometries. In 1D odd rings, a built-in domain wall yields reproducible, steplike reversal under a transverse field. In 2D finite samples, boundary-assisted switching and droplet growth are consistent with the observed evolution of the spin-structure factor. The highly coordinated hardware graphs suppress the ferromagnetic-like or-

dering, producing weak, near-minor loops. Overall, loop morphology is set by a small set of tunable parameters—geometry and coordination, the sweep protocol, and the ratio Γ/J —highlighting quantum annealers as controllable platforms for exploring nonequilibrium antiferromagnetic dynamics. The phenomena we quantify are reproduced on four different D-Wave quantum processors. There are intrinsic noise differences and timescale differences on the different analog quantum hardware, as well as different Hamiltonian energy scales. Despite these differences, the hysteretic characteristics we identify are robust across the different quantum processors.

Interestingly, some of the hysteresis experiments we present in this study share some qualities with measured magnetic hysteresis cycles in certain molecular magnets [60]. Similarities between these analog open-quantum system hysteresis experiments and molecular magnetic hysteresis deserves further investigation.

Future work can explore the dynamical origin of this hysteresis, its potential relationship with the qubit quantum dynamics, and also establish a closer connection with real magnetic materials and compare results from quantum annealers to those from laboratory experiments. Future studies should also investigate additional calibrated antiferromagnetic hysteresis cycle simulations (in particular of the couplers and of the flux bias offsets), in particular on the 2D systems where the antiferromagnetic order parameter data shows much larger statistical fluctuations.

This computational capability of near term analog quantum computers highlights that the full programmability of these devices inherently gives us the capability to examine phenomena which are hard to systematically study in magnetic laboratories—in this case, antiferromagnetic hysteresis.

ACKNOWLEDGMENTS

We thank Vivien Zapf for productive discussions. This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). The research presented in this article was supported by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under Project No. 20240032DR. This research used resources provided by the Los Alamos National Laboratory Institutional Computing Program. The authors would also like to thank the New Mexico Consortium, under subContract No. C2778, the Quantum Cloud Access Project (QCAP) for providing quantum computing resources. This is LANL Report No. LA-UR-25-28275.

DATA AVAILABILITY

The raw data are available on request to the corresponding author.

APPENDIX A: D-WAVE HARDWARE $A(s)$ AND $B(s)$ FUNCTIONS

Figure 11 plots the $A(s)$ and $B(s)$ schedules that define the D-Wave hardware Hamiltonian in Eq. (1), which allows us to then define physical ratios such as Γ/J . This is important specifically because all of the programmable units on the D-Wave QPUs are specified in terms of hardware-normalized parameters, and therefore it is necessary to extract the correct energy scales of the D-Wave anneal schedules to show where certain features of the magnetic hysteresis protocols appear in physical units, for example, at what applied field does magnetization reversal begin. Note that these frequencies are not in angular frequency units.

APPENDIX B: FLUX BIAS OFFSET BALANCING CALIBRATION

Due to the analog nature of the D-Wave quantum annealing hardware, it can be beneficial to calibrate the statistics of physical quantities that can be measured from the annealing simulations. In particular, analog errors and thermal effects can result in biased quantities. In this study we make use of a particular type of calibration procedure which balances the statistics obtained from the full-hardware-lattice-antiferromagnets. The update procedure used follows Ref. [55], and in particular is effectively a type of parallel-update steepest gradient descent machine learning process, and is very much an *autonomous* analog control correction procedure. This is reminiscent of gradient descent approaches in the training of neural networks [61,62]. The hardware control mechanism that is used here is the flux bias offset (FBO) control of each individual qubit [63] which can force the qubit to be biased towards spin up or down, and in this case we very slightly update the FBO proportionally to the size of the bias of the single-qubit magnetization (relative to the mean of the distribution across all qubits). The statistics balancing method in particular makes use of a correction towards the mean sampled observable (M_z) across all spins; in this case, we expect the net average magnetization to be zero, and the noncalibrated distributions also have a mean of approximately zero, which makes this correction method




FIG. 11. Overlaid $A(s)$ and $B(s)$ hardware specific functions for the four D-Wave QPUs used in this study. Note that the $A(s)$ functions on all four processors are effectively identical, which makes the lines visually indistinguishable.

work quite well in this case because its primary action is to reduce the variance of the measured single-spin magnetization on the lattice. The magnetization FBO update is relatively easy to perform because all of the qubits are in the same observable “orbit” [55], meaning we expect on average the magnetization of each qubit to be the same as all of the other qubits. Note, however, that this balancing did not adjust the J antiferromagnetic couplers; refining the programmed coupler energy values is possible to adjust as well, however boundaries in the lattice mean that different observable orbits must be constructed—which for the specific case of the highly complex hardware graphs would make the calibration significantly more computationally intensive in terms of total iterations and also parameter tuning. Therefore, we focused on FBO calibration, which yielded improved magnetization curve stability. This calibration procedure can in principle be applied to any programmed Ising model on analog hardware, but it has been especially useful for exploring sensitive phenomena, such as quantum quench dynamics and fluctuations near thermodynamic phase transitions. The FBO update rules applied to each qubit are uniform across the entire lattice, i.e., the bulk and edge qubits are treated identically, unlike in some prior studies [15].

Figure 12 shows the distribution of average magnetization per spin both before and after FBO statistics-balancing calibration. Figure 13 plots the values of the flux bias offsets, for every qubit in the hardware graph, as a function of update iteration, which in this case we terminate at a fixed iteration count of 3000 in order to make the simulations tractable in terms of compute time, and real-time backend usage. The flux bias offset constant step size used was $2e^{-6}$ (the actual update is proportional to the magnitude of the bias, and this is the constant by which the bias is multiplied by). This specific step size we found worked reasonably well from small scale empirical testing; larger stepsizes could in principle correct biases more quickly, but overcorrection and thereby numerical instability in the convergence was frequently observed for step sizes larger than $\approx 2e^{-6}$. The calibrated FBOs used in the subsequent hysteresis computations are simply the FBOs in the final learning iteration. Time drift is an important consideration in these D-Wave QPU computations (over long periods of time, D-Wave simulations do change measurably [48]), and therefore the complete set of hysteresis cycle data using the calibrated FBO values (shown in Fig. 10) were run within 48 hr of this calibration, with the aim of minimizing the effect of noise drift. At each update step, a total of 3000 anneal-readout cycles were performed (all in one device job call), and the subsequent update is performed using the (average) magnetization values for every site on the hardware lattice using all 3000 samples. Importantly, the anneal schedule used in this calibration procedure *was not* the hysteresis protocol—there was no anneal pause and no h -gain field. Instead, this calibration was performed only on the standard linear-ramp anneal schedule, using the final simulation time in the hysteresis protocol (in this case, always $11.2 \mu\text{s}$). This means the same FBO calibration (after the 3000 learning iterations), for each QPU, was used for each Γ/J and each “slice” of the longitudinal field ramps. More advanced calibration, which is also more QPU time intensive, could specifically calibrate each Γ/J , and each longitudinal field




FIG. 12. Single-site magnetization distributions, for each qubit in the hardware lattice, before flux bias offset calibration and after calibration for a full antiferromagnetic lattice defined on the entire hardware graph. 3000 gradient descent steps were used in total, and the flux bias offset step size is $2e^{-6}$ (with initial flux bias offset being uniformly zero). The standard deviation of the distribution of average single-site magnetization before (iteration 0) and after (iteration 3000) the FBO calibration is noted in the legend; in all cases we see a decrease in the variance of the single-site magnetization distribution, which is the goal of the statistical balancing process. Each iteration of this calibration process, on each QPU, uses exactly 3000 samples. The flux bias offset quantity calibration is shown in Fig. 13.

sweep step. Interestingly, Fig. 12 shows that some of the D-Wave processors are much easier to calibrate, presumably meaning they are less noisy, such as Advantage_system4.1, whereas Advantage_system7.1 appears to be the most noisy and also exhibited the least debiasing when this calibration protocol was applied as evidenced by having the widest magnetization distribution variance postcalibration. This is also indicated by the flux bias offset quantities as a function of iteration in Fig. 13, where we see that specifically in the case of Advantage_system7.1 and Advantage2_prototype2.6, some of the qubits did not converge in their FBO values, whereas Advantage_system4.1 and Advantage_system6.4 appear to have converged more quickly.

APPENDIX C: ADDITIONAL REPRESENTATIVE MAGNETIC SPIN STRUCTURE FACTORS

This appendix provides several additional sequences of averaged magnetic structure factors at various points along various 2D antiferromagnetic hysteresis simulations. Figures 14–19 show additional sequences, each obtained from different hysteresis cycles at various s values (equivalent to a physical Γ/J) and run on the several different D-Wave processors. In these structure factors, effects from the open boundary

conditions and finite lattice are prominent, particularly at full saturation. In the less saturated regions of the hysteresis cycles, the antiferromagnetic ordering becomes apparent with peaks at (π, π) in the square lattice. Due to the consistently applied longitudinal field, there is typically a strong peak at $(0, 0)$ which corresponds to ferromagnetic-like ordering, and becomes much stronger at saturation.

The Hamiltonian energy scale ratio Γ/J does change the characteristics of the magnetic ordering. Figure 15 (left) shows ordering at a low longitudinal field strength with $M_z \approx 0$, where the peak at $(0, 0)$ largely disappears. This peak of ordering disappears because Γ/J is strong in Fig. 15, whereas Fig. 16 displays magnetic structure factors also near magnetization $M_z \approx 0$, but here Γ/J is very weak, and this shows very different magnetic ordering compared to Fig. 15 with peaks of antiferromagnetic ordering along $q_x = \pm\pi, q_y \pm\pi$.

The general characterization we find is that qualitatively the magnetic ordering, as revealed by $|S(q)|$, is similar across the different physical experiment settings, including different QPUs. In other words, there are no strong outliers that would indicate a particular hardware issue, or an unexpected magnetic ordering. The computation of the structure factors is greatly accelerated by the Python libraries Numba and NumPy [64,65].




FIG. 13. Evolution of FBO values, for each single qubit in the hardware lattice, as the calibration process evolves up to 3000 iterations. The Ising model being sampled during this calibration is a full antiferromagnet defined lattice, using every coupler of each of the D-Wave hardware graphs. The initial flux bias offset starts at all zeros for all qubits, and each iteration has a maximum FBO step size of $2e^{-6}$.




FIG. 14. Magnetic structure factors $|S(q)|$ heatmaps (bottom row) within the 2D antiferromagnetic hysteresis cycles (averaged over 100 independent samples), run on `Advantage_system7.1` at $s = 0.7$. The Ising model being simulated is a 33×33 grid of (antiferromagnetically coupled) spins, with open boundary conditions. The top row shows the single hysteresis cycle in terms of M_z average magnetization, from which we extracted $|S(q)|$ heatmaps for four points along the first longitudinal field sweep, denoted by cyan asterisks (the order of the cyan asterisks have the same order as the $|S(q)|$ heatmaps). The dashed white line on the structure factor heat maps outlines the first Brillouin zone. The heat maps are log-scaled and scales are specific to each heatmap subplot. The order of the $|S(q)|$ heatmaps follows the same sweep direction denoted by the top hysteresis cycle plot (the SSF from the fully saturated regime is shown in the right hand subplot). Note the peaks at the corner of the Brillouin zone denoting large antiferromagnetic clusters coexist with a growing peak at the center denoting magnetization. The format of all other magnetic structure plots follow this same figure layout.




FIG. 15. Magnetic structure factors $|S(q)|$ heatmaps (bottom row) within the 2D antiferromagnetic hysteresis cycles, run on Advantage_system6.4 at $s = 0.65$.




FIG. 16. Magnetic structure factors $|S(q)|$ heatmaps (bottom row) within the 2D antiferromagnetic hysteresis cycles, run on Advantage_system6.4 at $s = 0.8$.




FIG. 17. Magnetic structure factors $|S(q)|$ heatmaps (bottom row) within the 2D antiferromagnetic hysteresis cycles, run on Advantage_system7.1 at $s = 0.6$.




FIG. 18. Magnetic structure factors $|S(q)|$ heatmaps (bottom row) within the 2D antiferromagnetic hysteresis cycles, run on Advantage_system4.1 at $s = 0.7$.




FIG. 19. Magnetic structure factors $|S(q)|$ heatmaps (bottom row) within the 2D antiferromagnetic hysteresis cycles, run on Advantage_system4.1 at $s = 0.65$.

APPENDIX D: EXPERIMENT STANDARD DEVIATION MAGNETIZATION MEASUREMENT QUANTITIES

The majority of the statistical uncertainty is due to the finite sampling effect (e.g., shot noise). However, because the protocol is fundamentally a sampling based protocol at intermediate slices of a continuous time computation, there is inherently some additional statistical noise during any magnetic reversal regions of the net magnetization plots. In Figs. 20 and 21 we show the standard deviation of several of the antiferromagnetic hysteresis cycles, which shows that overall the largest standard deviations are at most ≈ 0.035 . Comparing the standard deviation measures to the magnetization curves in Fig. 2 and Fig. 5, typically when Γ/J is large (near $s = 0.6$, $s = 0.7$) and the system undergoes full reversal, near applied field $H \approx 0$ the standard deviation peaks. This increase in variance at reversal is consistent with the stochastic effects of avalanche-mediated reversal and Barkhausen jumps.




FIG. 20. Standard deviation of the magnetization ($\sigma(M_z)$) measured observable from the hysteresis cycles reported in Fig. 2 of the 1D antiferromagnetic rings with periodic boundary conditions.




FIG. 21. Standard deviation of the magnetization ($\sigma(M_z)$) measured observable from the hysteresis cycles reported in Fig. 5 of 2D antiferromagnetic square lattices with open boundary conditions.

APPENDIX E: DOMAIN WALL DENSITY ON THE THREE OTHER D-WAVE QPUS

Figure 22 reports domain wall density during hysteresis cycles in the 1D (odd-spin ring) models from the remaining three QPUs, using the same plot style as Fig. 4.




FIG. 22. Domain wall density on the 1D antiferromagnetic rings during the hysteresis cycles as a function of applied longitudinal field H . Left-hand column shows the total, average, domain wall density, regardless of the sign of the domain wall. Right-hand column shows the spin-down ($\downarrow\downarrow$) domain wall density.

APPENDIX F: 2D LATTICE NÉEL ORDER PARAMETER QUANTITIES

We focus on the hysteretic region of s , and separate each hysteresis cycle by Γ/J in order to make visualization clear.

Figure 23 shows that the staggered antiferromagnetic order parameter is nonzero, and that it also forms hysteresis cycles. Data are from experiments on Advantage_system6.4. Figures 24–26 show the same set of data but from the other three quantum annealing processors. The occasional hysteresis is very distinct from the ferromagnetic-like magnetization curves seen in Fig. 5; for example, we see nonzero area between sweeps at $s \approx 0.75, 0.55$ in Fig. 23. Interestingly, at very small Γ/J we do see long-range antiferromagnetic ordering. This is expected because state transitions are much harder when the transverse field is weak. However, the data for the staggered Néel parameter are significantly noisier than average net magnetization in Fig. 5. This effect could in part originate in shot noise; also, the Néel order parameter could be more sensitive to hardware biases. In future studies we will apply statistics-balancing calibration refinement [55] on the local fields and on the couplers to reduce the noise. For the three devices where full magnetic M_z saturation was reached, M_s converges entirely to zero during the hysteresis cycle for the values of s corresponding to maximal hysteretic area in M_z (see Fig. 5).




FIG. 23. 2D hysteresis cycles in terms of the antiferromagnetic order parameter M_s . Data from Advantage_system6.4.




FIG. 24. 2D hysteresis cycles in terms of the antiferromagnetic order parameter M_s . Data from Advantage_system4.1.




FIG. 25. 2D hysteresis cycles in terms of the antiferromagnetic order parameter M_s . Data from Advantage_system7.1.




FIG. 26. 2D hysteresis cycles in terms of the antiferromagnetic order parameter M_s . Data from Advantage2_prototype2.6.

APPENDIX G: ADDITIONAL REPRESENTATIVE 2D SAMPLED CONFIGURATIONS DURING HYSTERESIS CYCLES

Figure 27 shows additional representative sampled configurations of the 2D antiferromagnetic model during a hysteresis cycle.




FIG. 27. Selected real-space spin configurations sampled on the Advantage_system7.1 QPU, illustrated as pixel plots where black pixels denote spin-down qubit measurements and cyan pixels denote spin-up qubit measurements (top row), from the middle of the magnetization reversal at various longitudinal field strengths of the first magnetic hysteresis sweep of a 33×33 2D square grid antiferromagnet at $s = 0.7$. Notice that the checkerboard antiferromagnetic ground state emerges in sections of the lattices. Also notice that the boundary effects from the open boundary conditions are noticeable in many of the measured spin configurations. Bottom row: Staggered domain height function τ from Eq. (13), corresponding to the same spin configurations on top, making the antiferromagnetic domains apparent.

APPENDIX H: D-WAVE QUANTUM ANNEALER HARDWARE EMBEDDINGS

All Ising models that were used in this study were hardware-native embeddings, meaning a spin-to-qubit mapping was uniquely defined for each model, for each D-Wave hardware graph. Figure 28 renders the 1D model embeddings, and Fig. 29 renders the 2D model embeddings. Note that the hardware graph defined antiferromagnetic models do not require an isomorphism embedding because they, by definition, use the entire hardware graph.




FIG. 28. 1D ring embeddings on the four hardware graphs of the D-Wave QPUs (Advantage2_prototype2.6 top left, Advantage_system4.1 top right, Advantage_system6.4 lower left, and Advantage_system7.1 lower right). Blue nodes and edges denote qubits and couplers used in the model embedding. Gray nodes and edges denote hardware not used in the embedding. Missing couplers or nodes from the hardware graph are inactive qubits/couplers that the user can not program (and in the case of Advantage_system7.1, the full Pegasus P_{16} graph was used for the visual rendering).




FIG. 29. 2D square grid, with open boundary conditions, embeddings on the four hardware graphs of the D-Wave QPUs (Advantage2_prototype2.6 top left, Advantage_system4.1 top right, Advantage_system6.4 lower left, and Advantage_system7.1 lower right). Blue nodes and edges denote qubits and couplers used in the model embedding. Gray nodes and edges denote hardware not used in the embedding. Missing couplers or nodes from the hardware graph are inactive qubits or couplers that the user cannot program (and in the case of Advantage_system7.1, the full Pegasus P_{16} graph was used for the visual rendering).

-
- [1] S. Morita and H. Nishimori, Mathematical foundation of quantum annealing, *J. Math. Phys.* **49**, 125210 (2008).
 - [2] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
 - [3] G. E. Santoro, R. Martonák, E. Tosatti, and R. Car, Theory of quantum annealing of an Ising spin glass, *Science* **295**, 2427 (2002).
 - [4] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
 - [5] A. D. King, C. D. Batista, J. Raymond, T. Lanting, I. Ozfidan, G. Poulin-Lamarre, H. Zhang, and M. H. Amin, Quantum annealing simulation of out-of-equilibrium magnetization in a spin-chain compound, *PRX Quantum* **2**, 030317 (2021).

- [6] R. Harris, Y. Sato, A. J. Berkley, M. Reis, F. Altomare, M. H. Amin, K. Boothby, P. Bunyk, C. Deng, C. Enderud, S. Huang, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, T. Lanting, R. Li, T. Medina, R. Molavi, R. Neufeld, *et al.*, Phase transitions in a programmable quantum spin glass simulator, *Science* **361**, 162 (2018).
- [7] S. N. Hearth, S. C. Morampudi, and C. R. Laumann, Quantum orders in the frustrated Ising model on the bathroom tile lattice, *Phys. Rev. B* **105**, 195101 (2022).
- [8] A. D. King, J. Carrasquilla, J. Raymond, I. Ozfidan, E. Andriyash, A. Berkley, M. Reis, T. Lanting, R. Harris, F. Altomare, K. Boothby, P. I. Bunyk, C. Enderud, A. Fréchet, E. Hoskinson, N. Ladizinsky, T. Oh, G. Poulin-Lamarre, C. Rich, Y. Sato, *et al.*, Observation of topological phenomena in a programmable lattice of 1,800 qubits, *Nature (London)* **560**, 456 (2018).
- [9] A. D. King, C. Nisoli, E. D. Dahl, G. Poulin-Lamarre, and A. Lopez-Bezanilla, Qubit spin ice, *Science* **373**, 576 (2021).
- [10] A. D. King, A. Nocera, M. M. Rams, J. Dziarmaga, R. Wiersema, W. Bernoudy, J. Raymond, N. Kaushal, N. Heinsdorf, R. Harris, K. Boothby, F. Altomare, M. Asad, A. J. Berkley, M. Boschnak, K. Chern, H. Christiani, S. Cibere, J. Connor, M. H. Dehn, *et al.*, Beyond-classical computation in quantum simulation, *Science* **388**, 199 (2025).
- [11] S. Zhou, D. Green, E. D. Dahl, and C. Chamon, Experimental realization of classical z_2 spin liquids in a programmable quantum device, *Phys. Rev. B* **104**, L081107 (2021).
- [12] P. Narasimhan, S. Humeniuk, A. Roy, and V. Drouin-Touchette, Simulating the transverse-field Ising model on the kagome lattice using a programmable quantum annealer, *Phys. Rev. B* **110**, 054432 (2024).
- [13] A. Lopez-Bezanilla, J. Raymond, K. Boothby, J. Carrasquilla, C. Nisoli, and A. D. King, Kagome qubit ice, *Nat. Commun.* **14**, 1105 (2023).
- [14] A. Lopez-Bezanilla, A. D. King, C. Nisoli, and A. Saxena, Quantum fluctuations drive nonmonotonic correlations in a qubit lattice, *Nat. Commun.* **15**, 589 (2024).
- [15] P. Kairys, A. D. King, I. Ozfidan, K. Boothby, J. Raymond, A. Banerjee, and T. S. Humble, Simulating the Shastry-Sutherland Ising model using quantum annealing, *PRX Quantum* **1**, 020320 (2020).
- [16] E. Pelofske, F. Barrows, P. Sathe, and C. Nisoli, Magnetic hysteresis experiments performed on quantum annealers, *Sci. Adv.* **12**, eaeb5192 (2026).
- [17] F. Barrows, E. Pelofske, P. Sathe, F. Caravelli, and C. Nisoli, Magnetic memory and hysteresis from quantum transitions in theory and experiments on quantum annealers, *Commun. Phys.* (2026).
- [18] S. A. Makhlof, F. T. Parker, and A. E. Berkowitz, Magnetic hysteresis anomalies in ferritin, *Phys. Rev. B* **55**, R14717 (1997).
- [19] C. Gilles, P. Bonville, H. Rakoto, J.-M. Broto, K. Wong, and S. Mann, Magnetic hysteresis and superantiferromagnetism in ferritin nanoparticles, *J. Magn. Magn. Mater.* **241**, 430 (2002).
- [20] S. Mijatović, S. Graovac, D. Spasojević, and B. Tadić, Hysteresis loop shape of field-driven antiferromagnetic-ferromagnetic bilayers in experimental conditions, *Phys. Rev. B* **111**, 064304 (2025).
- [21] M. J. Grzybowski, C. F. Schippers, O. Gomonay, K. Rubi, M. E. Bal, U. Zeitler, A. Koziół-Rachwał, M. Szpytma, W. Janus, B. Kurowska, S. Kret, M. Gryglas-Borysiewicz, B. Koopmans, and H. J. M. Swagten, Antiferromagnetic hysteresis above the spin-flop field, *Phys. Rev. B* **107**, L060403 (2023).
- [22] A. Paduan-Filho, C. C. Becerra, and F. Palacio, Hysteresis at the spin-flop transition in the antiferromagnets $K_2Fe(Cl_{1-x}Br_x)_5 \cdot H_2O$, *Phys. Rev. B* **43**, 11107 (1991).
- [23] L. Opherden, T. Bilitewski, J. Hornung, T. Herrmannsdörfer, A. Samartzis, A. T. M. N. Islam, V. K. Anand, B. Lake, R. Moessner, and J. Wosniza, Inverted hysteresis and negative remanence in a homogeneous antiferromagnet, *Phys. Rev. B* **98**, 180403(R) (2018).
- [24] B. K. Chakrabarti and M. Acharyya, Dynamic transitions and hysteresis, *Rev. Mod. Phys.* **71**, 847 (1999).
- [25] F. Broner, G. H. Goldsztein, and S. H. Strogatz, Dynamical hysteresis without static hysteresis: Scaling laws and asymptotic expansions, *SIAM J. Appl. Math.* **57**, 1163 (1997).
- [26] R. Moessner and S. L. Sondhi, Ising models of quantum frustration, *Phys. Rev. B* **63**, 224401 (2001).
- [27] H. Zou, Y. Cui, X. Wang, Z. Zhang, J. Yang, G. Xu, A. Okutani, M. Hagiwara, M. Matsuda, G. Wang, G. Mussardo, K. Hódsági, M. Kormos, Z. He, S. Kimura, R. Yu, W. Yu, J. Ma, and J. Wu, E_8 spectra of quasi-one-dimensional antiferromagnet $BaCo_2V_2O_8$ under transverse field, *Phys. Rev. Lett.* **127**, 077201 (2021).
- [28] Y. Cui, H. Zou, N. Xi, Z. He, Y. X. Yang, L. Shu, G. H. Zhang, Z. Hu, T. Chen, R. Yu, J. Wu, and W. Yu, Quantum criticality of the Ising-like screw chain antiferromagnet $SrCo_2V_2O_8$ in a transverse magnetic field, *Phys. Rev. Lett.* **123**, 067203 (2019).
- [29] C. Liu, C.-J. Huang, and G. Chen, Intrinsic quantum Ising model on a triangular lattice magnet $TmMgGaO_4$, *Phys. Rev. Res.* **2**, 043013 (2020).
- [30] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, *et al.*, Quantum annealing with manufactured spins, *Nature (London)* **473**, 194 (2011).
- [31] T. Lanting, A. J. Przybysz, A. Y. Smirnov, F. M. Spedalieri, M. H. Amin, A. J. Berkley, R. Harris, F. Altomare, S. Boixo, P. Bunyk, N. Dickson, C. Enderud, J. P. Hilton, E. Hoskinson, M. W. Johnson, E. Ladizinsky, N. Ladizinsky, R. Neufeld, T. Oh, I. Perminov, *et al.*, Entanglement in a quantum annealing processor, *Phys. Rev. X* **4**, 021041 (2014).
- [32] P. I. Bunyk, E. M. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, A. J. Przybysz, and J. Whittaker, Architectural considerations in the design of a superconducting quantum annealing processor, *IEEE Trans. Appl. Supercond.* **24**, 1 (2014).
- [33] N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, P. Bunyk, J. Cai, E. M. Chapple, P. Chavez, F. Cioata, T. Cirip, P. deBuen, M. Drew-Brook, C. Enderud, S. Gildert, F. Hamze, J. P. Hilton, E. Hoskinson, K. Karimi, *et al.*, Thermally assisted quantum annealing of a 16-qubit problem, *Nat. Commun.* **4**, 1903 (2013).
- [34] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M. C. Thom, E. M. Chapple, J. Wang, B. Wilson, *et al.*, Experimental investigation of an eight-qubit unit

- cell in a superconducting optimization processor, *Phys. Rev. B* **82**, 024511 (2010).
- [35] R. Harris, A. J. Berkley, M. W. Johnson, P. Bunyk, S. Govorkov, M. C. Thom, S. Uchaikin, A. B. Wilson, J. Chung, E. Holtham, J. D. Biamonte, A. Y. Smirnov, M. H. S. Amin, and A. M. van den Brink, Sign- and magnitude-tunable coupler for superconducting flux qubits, *Phys. Rev. Lett.* **98**, 177001 (2007).
- [36] M. W. Johnson, P. Bunyk, F. Maibaum, E. Tolkacheva, A. J. Berkley, E. M. Chapple, R. Harris, J. Johansson, T. Lanting, I. Perminov, E. Ladizinsky, T. Oh, and G. Rose, A scalable control system for a superconducting adiabatic quantum optimization processor, *Supercond. Sci. Technol.* **23**, 065004 (2010).
- [37] A. D. King, S. Suzuki, J. Raymond, A. Zucca, T. Lanting, F. Altomare, A. J. Berkley, S. Ejtemaee, E. Hoskinson, S. Huang, E. Ladizinsky, A. J. R. MacDonald, G. Marsden, T. Oh, G. Poulin-Lamarre, M. Reis, C. Rich, Y. Sato, J. D. Whittaker, J. Yao, *et al.*, Coherent quantum annealing in a programmable 2,000 qubit Ising chain, *Nat. Phys.* **18**, 1324 (2022).
- [38] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud, E. Hoskinson, S. Huang, E. Ladizinsky, A. J. R. MacDonald, G. Marsden, R. Molavi, T. Oh, G. Poulin-Lamarre, M. Reis, C. Rich, *et al.*, Quantum critical dynamics in a 5,000-qubit programmable spin glass, *Nature (London)* **617**, 61 (2023).
- [39] M. H. Amin, Searching for quantum speedup in quasistatic quantum annealers, *Phys. Rev. A* **92**, 052323 (2015).
- [40] Z. Morrell, M. Vuffray, A. Y. Lokhov, A. Bärttschi, T. Albash, and C. Coffrin, Signatures of open and noisy quantum systems in single-qubit quantum annealing, *Phys. Rev. Appl.* **19**, 034053 (2023).
- [41] J. Marshall, E. G. Rieffel, and I. Hen, Thermalization, freeze-out, and noise: Deciphering experimental quantum annealers, *Phys. Rev. Appl.* **8**, 064025 (2017).
- [42] L. Buffoni and M. Campisi, Thermodynamics of a quantum annealer, *Quantum Sci. Technol.* **5**, 035013 (2020).
- [43] N. Chancellor, S. Szoke, W. Vinci, G. Aeppli, and P. A. Warburton, Maximum-entropy inference with a programmable annealer, *Sci. Rep.* **6**, 22318 (2016).
- [44] J. Nelson, M. Vuffray, A. Y. Lokhov, and C. Coffrin, Single-qubit fidelity assessment of quantum annealing hardware, [arXiv:2104.03335](https://arxiv.org/abs/2104.03335).
- [45] C. Tüysüz, A. Jayakumar, C. Coffrin, M. Vuffray, and A. Y. Lokhov, Learning response functions of analog quantum computers: Analysis of neutral-atom and superconducting platforms, [arXiv:2503.12520](https://arxiv.org/abs/2503.12520).
- [46] M. Vuffray, C. Coffrin, Y. A. Kharkov, and A. Y. Lokhov, Programmable quantum annealers as noisy Gibbs samplers, *PRX Quantum* **3**, 020317 (2022).
- [47] T. Zaborniak and R. de Sousa, Benchmarking Hamiltonian noise in the D-Wave quantum annealer, *IEEE Trans. Quantum Eng.* **2**, 1 (2021).
- [48] E. Pelofske, G. Hahn, and H. N. Djidjev, Noise dynamics of quantum annealers: Estimating the effective noise using idle qubits, *Quantum Sci. Technol.* **8**, 035005 (2023).
- [49] T. Lanting, M. H. Amin, C. Baron, M. Babcock, J. Boschee, S. Boixo, V. N. Smelyanskiy, M. Foygel, and A. G. Petukhov, Probing environmental spin polarization with superconducting flux qubits, [arXiv:2003.14244](https://arxiv.org/abs/2003.14244).
- [50] C. McCreesh, P. Prosser, and J. Trimble, The Glasgow subgraph solver: Using constraint programming to tackle hard subgraph isomorphism problem variants, in *International Conference on Graph Transformation*, edited by F. Gadducci and T. Kehler, (Springer, Cham, 2020), pp. 316–324.
- [51] J. Cai, W. G. Macready, and A. Roy, A practical heuristic for finding graph minors, [arXiv:1406.2741](https://arxiv.org/abs/1406.2741).
- [52] N. Dattani, S. Szalay, and N. Chancellor, Pegasus: The second connectivity graph for large-scale quantum annealing hardware, [arXiv:1901.07636](https://arxiv.org/abs/1901.07636).
- [53] K. Boothby, P. Bunyk, J. Raymond, and A. Roy, Next-generation topology of D-Wave quantum processors, [arXiv:2003.00133](https://arxiv.org/abs/2003.00133).
- [54] K. Boothby, A. D. King, and J. Raymond, Zephyr topology of D-Wave quantum processors (2021), https://www.dwavesys.com/media/2uznec4s/14-1056a-a_zephyr_topology_of_d-wave_quantum_processors.pdf.
- [55] K. Chern, K. Boothby, J. Raymond, P. Farré, and A. D. King, Tutorial: Calibration refinement in quantum annealing, *Front. Comput. Sci.* **5**, 1238988 (2023).
- [56] A. Ali, H. Xu, W. Bernoudy, A. Nocera, A. D. King, and A. Banerjee, Quantum quench dynamics of geometrically frustrated Ising models, *Nat. Commun.* **15**, 10756 (2024).
- [57] P. Sathe, A. D. King, S. M. Mniszewski, C. Coffrin, C. Nisoli, and F. Caravelli, Classical criticality via quantum annealing, *Nat. Commun.* **17**, 856 (2026).
- [58] E. Pelofske, G. Hahn, and H. N. Djidjev, Reducing quantum annealing biases for solving the graph partitioning problem, in *Proceedings of the 18th ACM International Conference on Computing Frontiers (ACM, 2021)*, pp. 133–139.
- [59] K. Nishimura, H. Nishimori, and H. G. Katzgraber, Griffiths-McCoy singularity on the diluted Chimera graph: Monte Carlo simulations and experiments on quantum hardware, *Phys. Rev. A* **102**, 042403 (2020).
- [60] W. Jiang, V.-c. Lo, B.-D. Bai, and J. Yang, Magnetic hysteresis loops in molecular-based magnetic materials $\text{AFe}^{\text{II}}\text{Fe}^{\text{III}}(\text{C}_2\text{O}_4)_3$, *Phys. A* **389**, 2227 (2010).
- [61] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* **86**, 2278 (2002).
- [62] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Learning representations by back-propagating errors, *Nature (London)* **323**, 533 (1986).
- [63] R. Harris, T. Lanting, A. J. Berkley, J. Johansson, M. W. Johnson, P. Bunyk, E. Ladizinsky, N. Ladizinsky, T. Oh, and S. Han, Compound Josephson-junction coupler for flux qubits with minimal crosstalk, *Phys. Rev. B* **80**, 052506 (2009).
- [64] S. K. Lam, A. Pitrou, and S. Seibert, Numba: A LLVM-based Python JIT compiler, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC* (Association for Computing Machinery, New York, NY, USA, 2015).
- [65] C. R. Harris, K. Jarrod Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).