

Supplemental Material for “Resonant false vacuum decay in two dimensions on a 4000-qubit quantum annealer”

CONTENTS

S1. QPU embeddings	1
S2. Measurement procedure	1
S3. Estimating the time duration of the resonant evolution	3
S4. Longitudinal field sweep	3
S5. Magnetization and classical energy scaling	4
S6. Extended data for characterization of bubble expansion dynamics	6
S7. Optimization of the scaling collapse	13
S8. Alternate stochastic circuits	13
References	13

S1. QPU EMBEDDINGS

We generate an embedding of a two-dimensional torus on 4593 qubits of the Advantage2_system1.6 graph (Fig. S1) and 4579 qubits of the Advantage2_system1.13 graph (Fig. S2) by connecting individual layers of a two-dimensional square lattice. A single layer is constructed by connecting two vertical and two horizontal qubits positioned opposite one another within a single Zephyr cell and then repeating this pattern across all cells. The Zephyr graph architecture allows four independent layers to be constructed using the same Zephyr cells without interlayer couplings [1]. Periodic boundary conditions are implemented by connecting the edges of these layers such that the resulting logical system forms a twisted torus.

Some qubits and couplers in the full Zephyr graph are absent on the physical device due to fabrication defects. To minimize the impact of these defects on the simulations, and to create larger regions suitable for studying bubble expansion, the connectivity of qubits within defective cells is adjusted so that defects are clustered as much as possible within a single layer. As a result, the corresponding defects are localized within a restricted region of the logical system, allowing larger defect-free regions in the remaining parts of the lattice.

The Advantage2_system1.6 was used to perform simulations at constant h_z while sweeping over different values of h_z . During these simulations, qubits missing more than one coupling to neighboring spins (shown on the right panel of Fig. S1), were manually biased with a local field h_z using the flux bias offset parameter of the D-Wave device such that they did not participate in the dynamics. This was done to preserve the defect-free resonance conditions, since missing spins and couplers introduce additional effects. For example, in a defect-free system, a bubble of size 1 nucleates at $h_z = 4J$, corresponding to the point at which the energy cost of creating four domain walls equals the energy gain from flipping a spin. In the presence of defects, however, a size-1 bubble can nucleate adjacent to a defect already at $h_z = 3J$, since only three domain walls must be created. Similar effects occur for bubbles of larger sizes.

The Advantage2_system1.13 was used to perform simulations studying bubble expansion under a ramp protocol. In these simulations, qubits with two or more missing neighbouring couplings (shown on the middle panel of Fig. S2) were excluded from the dynamics using local h_z fields in order to isolate the spreading of the manually seeded bubble. Since the dynamics were studied at the resonant condition $h_z = 2J$, qubits with two missing couplings can themselves become resonant nucleation sites, as flipping such a spin requires the creation of only two domain walls. These flipped spins can then act as seeds for additional bubbles that spread through the resonance mechanism.

S2. MEASUREMENT PROCEDURE

The programmable D-Wave Hamiltonian is given by




FIG. S1. Logical square-lattice configurations with periodic boundary conditions realized on the Advantage2_system1.6 quantum annealer. (left) Black squares indicate missing logical spins arising from unavailable physical qubits. (middle) Orange squares denote logical spins with two or more missing neighboring logical connections due to non-functional physical couplers. (right) Orange squares denote logical spins with at least one missing connection to a neighboring spin resulting from non-functional physical couplers.




FIG. S2. Logical square-lattice configurations with periodic boundary conditions realised on the Advantage2_system1.13 quantum annealer. (left) Black squares indicate missing logical spins arising from unavailable physical qubits. (middle) Orange squares denote logical spins with two or more missing neighboring logical connections due to non-functional physical couplers. (right) Orange squares denote logical spins with at least one missing connection to a neighboring spin resulting from non-functional physical couplers.

$$\hat{\mathcal{H}}_{\text{DW}}(t) = -\frac{A(s)}{2} \sum_i \hat{\sigma}_i^x + \frac{B(s)}{2} \left(g(t) \sum_i h_i \hat{\sigma}_i^z + \sum_{i>j} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z \right), \quad (\text{S1})$$

where $\hat{\sigma}_i^\alpha$ are Pauli operators acting on qubit i , h_i are programmable longitudinal fields, and J_{ij} are programmable couplings between qubits i and j . The couplings are nonzero only for qubits that are physically connected in the quantum processing unit. The parameter $s = t/t_a \in [0, 1]$ is the normalized annealing coordinate, with $A(s)$ and $B(s)$ setting the energy scales of the transverse-field and Ising terms, respectively. The annealing schedule $s(t)$ is specified by a series of user-defined points (t_i, s_i) , between which the schedule is linearly interpolated. Similarly, the time-dependent gain function $g(t)$, defined by a set of control points (t_i, g_i) , provides control of the longitudinal fields through the combination $g(t)h_i$. To implement the mixed-field Ising model studied here, we restrict J_{ij} to nearest-neighbor ferromagnetic couplings on the logical lattice and choose a uniform coupling strength $J_{ij} = J_{\text{prog}} = -0.3$.

The longitudinal fields are set to be uniform, $h_i = h_{\text{prog}}$. Factoring out the overall scale $B(s)/2$, the Hamiltonian can then be mapped onto the mixed-field Ising form, with the effective transverse and longitudinal fields given by $h_x = \frac{A(s)}{B(s)|J_{\text{prog}}|}$, $h_z(t) = -g(t)h_{\text{prog}}/J_{\text{prog}}$. This mapping allows both h_x and h_z to be controlled independently during the annealing process through the annealing schedule $s(t)$ and the gain schedule $g(t)$, respectively.

To realize the dynamics, we employed a reverse-annealing protocol specified by the schedule $[(0, 1), (t_1, s_{h_x}), (t_2, s_{h_x}), (t_3, 1)]$, where s_{h_x} is chosen such that $h_x = \frac{A(s_{h_x})}{B(s_{h_x})|J_{\text{prog}}|}$. Starting from a classical spin configuration at $h_x = 0$, the system is reverse annealed to a point with a small but finite transverse field, held there for a pause time $\tau = t_2 - t_1$, and subsequently annealed back to $h_x = 0$ for readout. We read out the states of all the qubits in the computational basis. Typical values are $t_1 = 1.0 \mu\text{s}$, $\tau \in [0.25, 7.0] \mu\text{s}$, and $t_3 - t_2 = 200 \text{ ns}$.

Two different protocols for the longitudinal field were implemented using the device's h -gain feature. In the first protocol, $g(t)$ is held constant throughout the pause, corresponding to a fixed longitudinal field h_z . We repeat this procedure 100 times and average over repetitions and individual qubit measurements to get the average magnetization. Repeating the experiment for different values of h_z allows us to map out the resonance structure shown in Fig. S3. In the second protocol, $g(t)$ is linearly ramped from $g(t_1) = 1$ to $g(t_2) = -1$, with $h_{\text{prog}} = -0.7$, causing the longitudinal field to approach into the $|h_z| = 2J$ resonance while maintaining a finite transverse field. An example of the schedule is shown in Fig. S4. By varying the duration of the ramp, we control the time spent within the resonant region. This allows us to directly probe the resonant interface-growth dynamics and determine how the growth depends on the time near the resonance. The maximum duration of the ramp for all h_x values is set to 5×10^4 qubit time units which are determined by the energy scale as $2\hbar/(B(s_{h_x})|J_{\text{prog}}|)$.

To study the spreading dynamics in detail, we manually seed the initial bubble as a single flipped spin, which is fixed in place throughout the simulation using a local h_z field implemented via the flux-bias control parameter. This procedure ensures that the initial position of the bubble is well-defined and identical across realizations, allowing us to systematically study the subsequent time evolution of a single nucleated defect. In practice, small inhomogeneities between qubits can lead to biases toward particular bubble configurations. To mitigate these effects, we repeat the protocol using 10 different seed qubits as initial flipped spins. For each seed location, we perform 100 runs at each ramp duration. The resulting measurements are then averaged over all seed positions, thereby reducing qubit to qubit variations and improving statistical robustness. The locations of all seed spins used in the experiments are shown in Fig. S5.

S3. ESTIMATING THE TIME DURATION OF THE RESONANT EVOLUTION

To estimate the time spent by the system within the resonance regime, we first determine the resonance width at $h_x = 0.1J$ from the sweep shown in Fig. S3. The ramp protocol is designed to terminate at a final value of $h_z \approx 2.33J$, which approximately corresponds to the center of the resonance peak observed in the sweep. Although the longitudinal field remains near its resonant value during the final measurement quench of h_x , the transverse field is rapidly reduced, strongly suppressing the dynamics and effectively freezing the state of the system. We therefore neglect this contribution when estimating the time spent in resonance. A further complication is the dependence of the resonance width on h_x . However, since the width is estimated from measurements performed at the largest transverse field considered in the simulations, we can treat it as constant, which provides an upper bound on the resonance width and, consequently, on the time spent within the resonant region. The relevant timescale is therefore determined primarily by the duration of the h_z ramp. For ramp times ranging from $0.26 \mu\text{s}$ to $7 \mu\text{s}$, and noting that the protocol sweeps the field from $+h_z$ to $-h_z$, a simple algebra estimate yields a time within the resonance window of approximately 8 ns to 200 ns from the shortest ramp length to the longest.

S4. LONGITUDINAL FIELD SWEEP

The sequence of the various resonances is directly visible in the constant- h_z sweep shown in Fig. S3, where the measured magnetization M is plotted as a function of longitudinal field and evolution time. A detailed scan over h_z reveals the expected sequence of resonant features, with the lowest-order resonances appearing most prominently and the higher-order two- and three-spin resonances progressively weaker, consistent with their smaller tunneling amplitudes. While the observed resonance structure agrees qualitatively with theoretical predictions, the resonances occur at slightly shifted values of h_z . This discrepancy is most likely due to small differences between the nominal programmed field and the actual longitudinal field realized in the physical device. The effect is further amplified by the use of a slightly reduced programmed coupling strength J_{prog} , which increases the sensitivity of the resonance positions to such calibration offsets. To verify that the observed dynamics correspond to the predicted resonant




FIG. S3. (Left) Experimental sweep of the longitudinal field on the annealer, showing the magnetization M as a function of h_z and evolution time t for the constant- $h_z(t)$ reverse-annealing protocol. (Right) Representative spin configurations and corresponding bubble-size distributions $n(s)$ measured at selected resonance points.

bubble nucleation processes, we analyze both representative spin configurations and the statistics of domain sizes at the resonance points. Specifically, we consider the number density $n(s)$ of connected domains (bubbles) of the true-vacuum phase containing s spins, defined as $n(s) = \frac{N(s)}{N_{\text{sys}} N_{\text{reps}}}$, where $N(s)$ is the number of bubbles of size s , N_{sys} is the total number of spins in the system, and N_{reps} is the number of snapshots included in the analysis. Representative snapshots together with the corresponding bubble-size distributions are shown on the right-hand side of the figure. These measurements confirm the expected resonance dynamics: bubbles of the resonant size dominate the distribution at each resonance, while larger bubbles appear with progressively lower density as a consequence of decoherence-driven growth following nucleation. The one-spin resonance is intentionally measured at a slightly offset value of h_z , since its dynamics are fast enough that, even at the shortest accessible evolution times, the system has already largely decayed to the true-vacuum ground state through a combination of single-spin bubble nucleation and subsequent noise-assisted relaxation. This effect is particularly pronounced because the sweep was performed at $h_x = 0.1J$, a value chosen to make even the higher-order resonances experimentally observable. The snapshots further confirm the nature of the individual resonant processes, showing isolated one-spin nucleation near $|h_z| = 4J$, small two- and three-spin bubbles at lower fields, and the onset of the $2J$ resonance where bubble growth becomes energetically allowed. Taken together, the magnetization data, real-space spin configurations, and bubble-size statistics provide direct experimental evidence that the annealer resolves the predicted sequence of resonant bubble nucleation and growth processes.

S5. MAGNETIZATION AND CLASSICAL ENERGY SCALING

The observed dependence of the dynamics on h_x reveals a nontrivial scaling behavior that can be understood from the underlying quantum dynamics. For a system held at a fixed resonant field, the leading-order growth process is governed by an effective Hamiltonian whose dynamics depend on the combination $h_x t$. In our protocol, however, the seeded bubble evolves while the longitudinal field $h_z(t)$ is swept through the resonance. In this case, both the experimental data and numerical simulations indicate an effective scaling with $h_x^2 t$ as shown in Fig. S6. This behavior arises because the width of the resonance is itself proportional to h_x . During a linear sweep of h_z , the duration for which the system remains within the resonant region therefore scales as $h_x t$. Since the resonant growth rate is also




FIG. S4. Ramp protocol used to characterize bubble expansion dynamics. A configuration containing a single initial seed spin is prepared during a time interval $t_1 = 1 \mu\text{s}$ by slowly ramping h_x from zero to a small target value while maintaining a constant positive h_z . Subsequently, h_z is ramped to the opposite sign over a timescale τ . Finally, h_x is reduced to zero as rapidly as possible during the interval $t_3 - t_2$, after which the final spin configuration is measured. The protocol is repeated 100 times for each value of τ and for each of the 10 initial seed-spin locations, resulting in a total of 1000 measurements for each data point.




FIG. S5. Configuration showing the spin locations used as the initial bubble seeds in the simulations characterizing bubble expansion. Black squares indicate spins absent from the system due to defects, while green squares denote the positions of the initial seeded bubbles.

proportional to h_x , the total amount of growth accumulated during the sweep is expected to depend on the product of these two factors, leading to an overall scaling in the variable $h_x^2 t$. Fig. S6 compares the experimental results with tree-tensor-network simulations performed using a comparable ramp protocol. In the numerical calculations, the ramp is chosen such that it terminates exactly at the resonance, closely matching the experimental procedure. The magnetization is shown both as a function of the ramp duration τ and as a function of the rescaled variable $h_x^2 \tau$. While the data obtained at different values of h_x exhibit distinct behavior when plotted against τ , they collapse onto a common curve when plotted against $h_x^2 \tau$. The same collapse is observed in the tensor-network simulations, indicating that the dynamics in both cases are controlled by the scaling variable $h_x^2 \tau$.

To further probe the coherence of the dynamics, we analyze two quantities extracted from the final spin configurations. The first is the classical Ising energy $E = h_z \sum_i s_i^z - \sum_{\langle ij \rangle} s_i^z s_j^z$, where $s_i = \pm 1$ denotes the measured state of spin i . The second is the number of disconnected domains of the true-vacuum phase, obtained by counting all connected clusters of flipped spins. Both quantities are averaged over all snapshots and over the ten different seed-spin locations used in the experiment. The resulting data are shown in Fig. S7. At short times, the classical energy remains approximately conserved, indicating that the dynamics remains mostly driven by the resonance condition rather than being dominated by generic off-resonant relaxation processes. This behavior is consistent with coherent resonant growth and provides evidence that the annealer is accessing the intended quantum dynamics. At longer




FIG. S6. Magnetization measured after the ramp protocol in tensor networks numerical simulations (TN) (top) and in quantum simulations (QS) (bottom) as function of ramp duration τ (left) and as a function a rescaled time axis $h_x^2 \tau$ (right), plotted for different h_x values. The rescaled curves for both methods collapse onto a single curve.

times, deviations from energy conservation become apparent. Initially, these manifest as an increase in the measured classical energy, which coincides with an increase in the number of domains. This behavior can be attributed to noise-induced spin flips that create additional domain walls and hence raise the classical energy. The resulting small bubbles then act as new seed domains, which can themselves undergo resonant expansion. At still later times, the classical energy begins to decrease while the number of domains remains approximately constant. This behavior is consistent with incoherent relaxation toward the true-vacuum ground state, as decoherence progressively fills in the bulk of the already expanded bubbles. A comparison of the scaling behavior of these observables reveals a further connection to the resonant growth mechanism. The classical energy exhibits the same $h_x^2 t$ scaling observed in the magnetization dynamics, as evidenced by the collapse of energy curves obtained at different values of h_x onto a single universal trajectory when plotted against $h_x^2 t$. In contrast, the number of domains displays a somewhat slower scaling behavior, suggesting that domain nucleation and domain growth are governed by partially distinct processes.

S6. EXTENDED DATA FOR CHARACTERIZATION OF BUBBLE EXPANSION DYNAMICS

The complete dataset for the magnetization, presented both as two-dimensional arrays and as horizontal line cuts through the center of the system, is presented for each initial seed-spin location in Figs. S8 and S9. To reduce the effects of local inhomogeneities and defects, the data from all seed-spin locations are combined. The resulting averaged data are presented in Fig. S10 together with the corresponding kink-density calculations, defined as $\hat{K}_{i,i+1} = \frac{1}{2} \sum_i (1 - \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z)$. To characterize the propagation dynamics more quantitatively, we evaluate the equal-time correlation function

$$\langle \hat{\sigma}_{-j}^z \hat{\sigma}_j^z \rangle$$

and the connected correlation function

$$\langle \hat{\sigma}_{-j}^z \hat{\sigma}_j^z \rangle_c = \langle \hat{\sigma}_{-j}^z \hat{\sigma}_j^z \rangle - \langle \hat{\sigma}_{-j}^z \rangle \langle \hat{\sigma}_j^z \rangle,$$

where $\hat{\sigma}_j^z$ denotes the magnetization of a spin located a horizontal distance j from the center of the initially seeded bubble.




FIG. S7. Classical energy (top) and number of domains of the true vacuum state (bottom) as a function of τ , plotted for several h_x values (left) and as a function of a rescaled time axis $h_x^2\tau$ (right). Rescaled energy curves collapse onto a single trajectory. Contrastingly, domain number curves exhibit slower dynamics.

To extract the position of the bubble boundary, threshold values are chosen for both correlation functions. The boundary position is then defined by the distance j and evolution time τ at which the corresponding threshold is reached. The thresholds used in the analysis are 0.93 for the equal-time correlation function and 5.0×10^{-5} for the connected correlator.

The nature of the bubble expansion is quantified by fitting the extracted boundary positions to the power-law form

$$r(t) = a t^p + b,$$

where a , b , and p are fitting parameters. The offset b accounts for the fact that the shortest experimental ramp times do not correspond exactly to $t = 0$. The uncertainty in the fitted exponent p is obtained from the covariance matrix returned by a weighted nonlinear least-squares fit. The boundary-position uncertainties are estimated from the spread of the extracted bubble front locations obtained for different initial seed spin positions and are used as weights in the fit. The reported error on p is given by the square root of the corresponding diagonal element of the covariance matrix. The complete dataset of correlation plots and fits is shown in Fig. S11.

The fitted propagation exponent stays close to $p = 1$ over the explored range of h_x , indicating approximately ballistic growth of the resonant bubble within the experimentally measured timescales. At small values of h_x , the bubble expands only over a small number of spins, resulting in larger uncertainties in the extracted boundary positions and consequently less stable fits. At larger values of h_x , the bubble reaches the boundaries of the finite system and wraps around through the periodic boundary conditions. This leads to self-interference effects that artificially increase the apparent propagation velocity and result in fitted exponents exceeding $p = 1$.

To mitigate these effects, the fitting procedure is restricted to evolution times preceding the onset of significant self-interference. The fitting ranges were determined from the kink-density maps shown in Fig. S10, where the onset of boundary-induced effects can be identified. The corresponding truncation points are indicated by the endpoints of the bubble-boundary trajectories displayed in the correlation-function plots.

Furthermore, we provide the complete dataset of magnetization profiles obtained from the quantum simulations along a one-dimensional cut through the system for multiple evolution times and for all measured values of h_x in Fig. S12. The profiles are analyzed in both the comoving coordinate $(j/\tau - v)$ and the KPZ-rescaled coordinate $(j/\tau - v)\tau^{2/3}$, separately for the left and right sides of the initially seeded bubble.




FIG. S8. Time evolution of the measured local magnetization profile $-\langle\sigma_j^z\rangle$ along a horizontal cut through the center of the two-dimensional system, shown as a function of distance j from the initially prepared central seed and evolution time τ , for all initial seed locations and all measured h_x values.




FIG. S9. Magnetization profiles $-\langle\sigma_j^z\rangle$ at different times τ along a horizontal cut through the center of the two-dimensional system, shown as a function of distance j from the initially prepared central seed and evolution time τ , for all initial seed locations and all measured h_x values.




FIG. S10. Raw magnetization profiles $-\langle \hat{\sigma}_j^z \rangle$ (left), its time evolution (middle) and nearest-neighbor kink density $\langle \hat{K}_{j,j+1} \rangle$ (right) along a horizontal cut through the center of the two-dimensional system, shown as a function of distance j from the initially prepared central seed and evolution time τ averaged over all initial seed locations and shown for all measured h_x values.




FIG. S11. Equal-time correlation function $\langle \hat{\sigma}_{-j}^z \hat{\sigma}_j^z \rangle$ (left) and connected equal-time correlator $\langle \hat{\sigma}_{-j}^z \hat{\sigma}_j^z \rangle_c$ (right), where $\hat{\sigma}_j^z$ denotes the single-site magnetization at horizontal distance j from the central initialized bubble site. The black points mark the extracted edge positions of the correlation front, and the green line shows a power-law fit $c \cdot \tau^p + a$. Plotted for all measured h_x values.




FIG. S12. Magnetization profiles along a one-dimensional cut for several evolution times plotted in the comoving coordinate $(j/\tau - v)$ and in the KPZ-scaled coordinate $(j/\tau - v)\tau^{2/3}$ for both the left and the right side from the initial seed location. Plotted for all measured h_x values.

S7. OPTIMIZATION OF THE SCALING COLLAPSE

To further quantify the scaling of the wavefront broadening, we employed a variance-based collapse metric to determine the optimal scaling exponent numerically. Assuming a power-law scaling form, the spatial coordinate was rescaled according to $j' = \frac{j-v\tau}{\tau^\alpha}$, where α is the scaling exponent and v is an effective propagation velocity. For a given parameter α , all magnetization profiles measured at different times τ for a fixed value of h_x were interpolated onto a common grid in the rescaled coordinate j' . Denoting the interpolated profiles by $m_i(j')$, the quality of the collapse was quantified by the mean variance across all profiles,

$$L(\alpha) = \frac{1}{N_k} \sum_{k=1}^{N_k} \text{Var}_i[m_i(j'_k)], \quad (\text{S2})$$

where j'_k are the points of the common interpolation grid, $N_k = 100$ is the number of grid points, and Var_i denotes the variance over the different times τ_i . Explicitly,

$$\text{Var}_i[m_i(j'_k)] = \frac{1}{N_\tau} \sum_{i=1}^{N_\tau} [m_i(j'_k) - \bar{m}(j'_k)]^2,$$

with $\bar{m}(j'_k) = \frac{1}{N_\tau} \sum_{i=1}^{N_\tau} m_i(j'_k)$, where N_τ is the number of time points included in the analysis. The optimal scaling parameters were then obtained by numerically minimizing the cost function $L(\alpha)$.

The results of this analysis are shown in Fig. S13. For both sides of the seeded spin, we plot the collapse cost function as a function of the scaling exponent α , together with the corresponding magnetization profiles displayed in the optimally rescaled coordinate. The value $\alpha = 2/3$, corresponding to the KPZ universality class, is indicated for reference. Overall, the extracted exponents are broadly consistent with KPZ scaling. This agreement is particularly evident on the right-hand side of the seeded spin, where the cost function exhibits a well-defined minimum close to $\alpha = 2/3$ and the rescaled wavefronts collapse onto a common curve. The left-hand side shows somewhat weaker agreement, although the cost function remains relatively flat over a broad range of exponents, indicating that the data are not strongly sensitive to the precise value of α . Nevertheless, the optimal exponent remains compatible with the KPZ scaling within the uncertainty of the collapse procedure.

S8. ALTERNATE STOCHASTIC CIRCUITS

In the main text, we use a classically simulable stochastic circuit to study the long-time behavior of the system. In this Section, we explore variants of this circuit and show that it leads to the same ballistic growth in all cases.

First, Fig. S14(a) shows that changing the flip probability p only leads to a rescaling of time as long as $p \ll 1$.

We also explore a non-brickwork circuit in which all sites (both even and odd) can be updated at the same step. In order to avoid violating the constraint, we need to resolve “conflicts” when two neighboring sites should be flipped at the same time. Finding the minimum number of flips to remove to resolve all conflicts is equivalent to finding the minimum vertex cover of a graph where each vertex is a site that should flip and edges represent conflicts between them. This is NP-hard in general. However, for bipartite graphs the minimum vertex cover can be computed in polynomial time. Indeed, one can first find a maximum cardinality matching using Hopcroft–Karp algorithm [2] and then use König’s theorem to go from the matching to a vertex cover. Since the graph on which conflicts are to be resolved is a subgraph of the physical square lattice, it is bipartite and we can use this approach.

Since the bipartiteness of the graph is what motivated the brickwork approach in the first place, we also explore another method that is agnostic to it. For that, we use an heuristic algorithm to find the vertex cover that does not rely on the graph being bipartite. In practice, we turn to a local-ratio algorithm which, while usually not giving the global minimum vertex cover, is guaranteed to return a set of vertices with at most twice its cardinality [3]. All graph algorithms used relied on their implementation in the python package networkX [4]. The results for both of these variants are shown in Fig. S14(b).

Overall, we find that all approaches lead to similar results, emphasizing that the scaling does not depend on microscopic detail but simply on the form of the constraint.

[1] Boothby, K. and King A. D. and Raymond, J., *Zephyr Topology of D-Wave Quantum Processors*, Tech. Rep. (D-Wave Systems Inc., Burnaby, BC, Canada, 2021).




FIG. S13. Cost function metric values and magnetization profiles along a one-dimensional cut for several evolution times plotted in the rescaled coordinate $(j/\tau - v)\tau^\alpha$, where α represents the optimal scaling parameter determined using numerical optimization. Plotted for both the left and the right side from the initial seed location and all measured h_x values.




FIG. S14. Comparison between the different microscopic implementations of the stochastic circuit. (a) Brickwork circuit layout alternating between the update of even and odd sites, averaged over 300 trajectories. Up to a rescaling of time, changing the flip probability p has virtually no effect. (b) Comparison between the brickwork circuit (red, 3000 trajectories) and the alternative approaches where all sites can be updated together and where conflicts are resolved using either the bipartite structure of the graph (orange, 250 trajectories) or a heuristic method (green, 250 trajectories), all with $p = 0.1$. In all cases a linear growth of r is recovered, albeit with a slightly different slope. The all-update method leads to growth slower by a factor of approximately 1.23.

- [2] J. E. Hopcroft and R. M. Karp, An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs, *SIAM Journal on Computing* **2**, 225 (1973).
- [3] R. Bar-Yehuda and S. Even, A local-ratio theorem for approximating the weighted vertex cover problem, in *North-Holland Mathematics Studies*, Vol. 109 (Elsevier, 1985) pp. 27–45.
- [4] A. A. Hagberg, D. A. Schult, and P. J. Swart, Exploring network structure, dynamics, and function using networkx, in *Proceedings of the 7th Python in Science Conference*, edited by G. Varoquaux, T. Vaught, and J. Millman (Pasadena, CA USA, 2008) pp. 11 – 15.