

The Power of an Adversary in Glauber Dynamics

Byron Chin
Ankur Moitra
Elchanan Mossel

Department of Mathematics, Massachusetts Institute of Technology

Colin Sandon
Chair of Mathematical Data Science, EPFL

BYRONC@MIT.EDU

MOITRA@MIT.EDU

ELMOS@MIT.EDU

COLIN.SANDON@EPFL.CH

Editors: Shipra Agrawal and Aaron Roth

Abstract

Glauber dynamics are a natural model of dynamics of dependent systems. While originally introduced in statistical physics, they have found important applications in the study of social networks, computer vision and other domains. In this work, we introduce a model of corrupted Glauber dynamics whereby instead of updating according to the prescribed conditional probabilities, some of the vertices and their updates are controlled by an adversary. We study the effect of such corruptions on global features of the system. Among the questions we study are: How many nodes need to be controlled in order to change the average statistics of the system in polynomial time? And how many nodes are needed to obstruct approximate convergence of the dynamics? Given a specific budget, how can the adversary choose nodes to control to maximize the overall effect?

Our results can be viewed as studying the robustness of classical sampling methods and are thus related to robust inference. The proofs connect to classical theory of Glauber dynamics from statistical physics.

1. Introduction

Many models of great interest in mathematics have been motivated by the goal of explaining complex systems and their behavior. The phase transitions of magnets, the evolution of opinions on a social network, the spread of infections among a population, and the inter-dependencies between nodes of a biological network are important examples of systems of interest.

Statistical physics studies the macroscopic behaviors of large systems that emerge. Often times these systems are built from simple building blocks. Major efforts have been invested to describe the properties of such systems in equilibrium, and how they change under perturbations. For example, what is the effect of imposing boundary conditions? And how do small differences in the initialization of the system play out? The seminal works of [Martinelli \(1999\)](#); [Martinelli and Toninelli \(2010\)](#) study the relaxation to equilibrium with various boundary conditions imposed on the Ising model. The study of differences in initialization is often called “damage spreading” as coined by [Kauffman \(1969\)](#). Many works have performed physical experiments and simulations to estimate the impact of small changes in initialization on macroscopic properties in equilibrium, see e.g. [Stanley et al. \(1987\)](#). See [Puzzo and Albano \(2008\)](#) for a survey of the work and results in this area.

But what can we say about the effect of perturbations over short periods of time as the system evolves toward equilibrium? A natural question to ask in this framework is: what macroscopic behavior can an adversary *efficiently* induce in a physical system with control over some subset of its microscopic constituents? By the term *efficiently* we mean that we are interested in the effects

after at most a polynomial number of time steps of running the natural dynamics. We emphasize that our notion of an adversary is new and flexible, as it allows the nodes to act, collude, and change over time arbitrarily.

In this work, we focus on systems represented by an Ising model. The Ising model has wide-ranging applications, including as a model for opinion dynamics in a social network in [Liu et al. \(2010\)](#), as a model for computer networks in [Antonio et al. \(2010\)](#), and as a model for a biological system in [Majewski et al. \(2001\)](#). See [Lipowski \(2022\)](#) for an extensive collection of works studying the Ising Model from various perspectives. Analogous questions are similarly interesting in other settings such as the Potts model, other Markov random fields and spin glasses. In such settings the macroscopic behavior can be much more complex, and the effectiveness of the adversary becomes even more apparent.

The Glauber dynamics introduced by [Glauber \(1963\)](#) is the natural formalism by which to add a time component to the Ising model to describe its evolution and progress towards equilibrium. Indeed, prior work by [Montanari and Saberi \(2010\)](#) and [Lara et al. \(2019\)](#) used it to study the spread of innovation and disease across a network. More recently [Baldassarri et al. \(2022\)](#) studied the evolution of opinions under Glauber-type dynamics. Another interesting application is [Jiang et al. \(2012\)](#), where the Glauber dynamics are used to produce a scheduling algorithm on a wireless network. Each of these examples can be subject to adversarial behavior of collections of nodes or agents.

One setting to highlight is that of a social network. Each vertex of a graph represents a person, with the corresponding label being their opinion on a specified topic. The edges between vertices correspond to personal connections that may influence the opinions of both people. Encoding opinions in this way allows for the analysis of both the distribution of opinions over the entire population, as well as how it changes over time. The adversary in this case can be interpreted as an advertisement campaign or an external party seeking to influence the global population as efficiently as possible.

1.1. Our Contributions

We introduce a general model of adversarially perturbed dynamics called *corrupted Glauber dynamics*. And we establish the first baselines, in this natural model, of the power of adversarial corruptions.

For context, in the high temperature regime, it is known that the presence of boundary conditions does not alter the equilibrium measure by much. For example, the notion of strong spatial mixing introduced by [Martinelli and Olivieri \(1994\)](#) stipulates that in equilibrium, the correlation between spins decreases exponentially in the graph distance between the corresponding vertex sets. While in physical systems it is natural to assume that the boundary conditions are “far away”, in our setting, the set of controlled vertices can be distributed in such a way that every node is close to one or more of the corrupted vertices. Our goal is to study the effect of controlling general subsets and we show that controlling a small set of corrupted vertices has a small effect on the system as a whole.

Second, in the low temperature regime, it was proven in many settings, starting with [Peierls \(1936\)](#) that the equilibrium measure consists of two distinct phases, a positive phase and a negative phase. Moreover, we know that the intermediate configurations between these phases are super-polynomially unlikely under the equilibrium measure for many families of graphs, see for example [Levin and Peres \(2017\)](#). In this strong correlation regime, fixing even a single vertex to $+1$ can

heavily bias the configuration toward the positive phase at equilibrium. Surprisingly, this property does not always carryover when we restrict to a polynomial number of steps of the dynamics. In particular, we establish that control of a small constant fraction of the vertices is insufficient to flip the phase in expander graphs. On graphs that are less expanding such as the $n \times n$ grid, we establish results showing to change the phase in polynomial time requires $\Theta(n)$ vertices.

We also move away from the setting of fixed boundary conditions and take advantage of the flexibility of our definition to accommodate more complex goal. In particular, suppose the corrupted set of vertices behaves adversarially in more sophisticated and coordinated ways. In the low temperature regime, it is known that configurations away from one of the two equilibrium phases are super-polynomially unlikely. We consider how much control the adversary requires to force the dynamic configuration to remain in these unlikely states, which are most polarized, for super-polynomial time. We prove that in the mean field setting with n vertices, it suffices to control slightly more than \sqrt{n} vertices to induce polarization in the spin configuration for a super-polynomial amount of time. Thus a coordinated adversary can force a system to remain in a fragile state for an extended period of time.

Finally, we consider the optimization problem from the perspective of the adversary. Following up on how large of an effect the adversary can have, it is natural to ask if the adversary can efficiently determine which set of vertices produces this maximal effect. This question is closely related to the problem of influence maximization, initially studied by [Kempe et al. \(2003\)](#); [Mossel and Roch \(2010\)](#). We initiate the corresponding study of influence maximization in the short-term evolution of the model by answering some preliminary questions.

1.2. Main Results

The main purpose of our paper will be to analyze a new process, which we call the corrupted Glauber dynamics, in which some vertices update adversarially and in a coordinated fashion.

Definition 1 (Corrupted Glauber Dynamics) *Consider a graph G and a subset of vertices $A \subset V(G)$. The A -corrupted Glauber dynamics is a Markov Chain that follows the transitions of the standard Glauber dynamics, except when a vertex $v \in A$ is chosen to be updated, the update is over-ridden by a rule specified by an adversary. The rule can be any function of the current and past states, and has no requirement of computational efficiency. The adversary can update any of the corrupted vertices at any time. We use m -corrupted Glauber dynamics to refer to A -corrupted Glauber dynamics where A is unspecified other than $|A| = m$.*

Remark 2 *In the case of a particular set of vertices being pinned to a fixed state by the adversary, the resulting measure can be interpreted as the Ising model with boundary conditions imposed. See [Martinelli and Toninelli \(2010\)](#); [Dyer et al. \(2004\)](#); [Weitz \(2005\)](#) for a definition and discussion of boundary conditions.*

The Ising model will be the vehicle we use to study and quantify macroscopic changes that can be induced when certain sets of vertices are controlled by an adversary.

Definition 3 (Ising Model) *Let $G = (V(G), E(G))$ be a graph and $\Omega = \{+1, -1\}^{V(G)}$. The Ising Model with parameter β is defined by the probability measure μ on Ω where for $\sigma = (\sigma(v))_{v \in V(G)}$,*

$$\mu(\sigma) = \frac{e^{\beta H(\sigma)}}{Z(\beta)}$$

where $H(\sigma) = \sum_{uv \in E(G)} \sigma(u)\sigma(v)$ and $Z(\beta) = \sum_{\sigma \in \Omega} e^{\beta H(\sigma)}$ is the normalization constant. We refer to the ± 1 label assigned to a vertex v as its spin.

The parameter β is known as the inverse temperature, which governs the strength of the dependencies between spins on adjacent vertices in G . The magnetization is the canonical measure of the global state of an Ising model as it measures the average behavior of the system. In much of the discussion below we will assume that the adversary wishes to shift the magnetization in a certain direction.

Definition 4 (Magnetization) *Given a configuration $\sigma \in \Omega$ the magnetization $m(\sigma) := \sum_{v \in V(G)} \sigma(v)$ is the sum of all the spins over all the vertices.*

Our main results can be summarized in the following theorems. In the high temperature regime, it is known that boundary conditions, i.e. a small set of vertices that are far away from most vertices, induces a negligible change in the resulting equilibrium measure as measured by magnetization. One sense in which this is captured is the property of strong spatial mixing and correlation decay. In our models, it is not natural to consider only far away vertices, so we consider the more general question of the effect of a small vertex set. The following theorem shows that the two configurations can be coupled such that for a polynomial number of time steps, the number of vertices that differ is very likely to be small. In particular, controlling a small number of vertices cannot change the overall statistics by much in the high temperature setting.

Theorem 5 *Let G be an n -vertex, d -regular graph. For $\beta(d)$ sufficiently small, when controlling ϵn of the vertices, the probability that the magnetization changes by 2ϵ in polynomial time is exponentially small in n . In fact, the trajectory can be coupled with a standard dynamics such that difference in magnetization is at most 2ϵ for exponential time with high probability.*

Note that the theorem above implies both that in polynomial time it is impossible to deviate by more than 2ϵ and also that the same is true for the stationary measure. See Theorem 13 for a formal statement. In contrast, in the regime of low temperature, it is known that the equilibrium exhibits a non-zero magnetization Peierls (1936). As a result imposing homogeneous boundary conditions (as in Martinelli and Toninelli (2010)), or even fixing a single vertex, will select a phase of the equilibrium measure. Thus, it is natural to wonder if a small corrupted set can change the phase of the configuration efficiently. On the contrary, we prove that on an expander graph, a small fraction of the vertices does not suffice to dynamically flip the magnetization in a polynomial number of time steps. Thus there are important differences in the efficacy of adversarial corruptions that vary greatly depending on the inherent robustness and error correction properties of the underlying network.

Theorem 6 *Let G be an n -vertex d -regular α -edge expander graph. For $\beta(d, \alpha)$ sufficiently large and $\epsilon(d, \alpha)$ sufficiently small, when controlling ϵn vertices the stationary measure still consists of two phases – one with positive magnetization and one with negative magnetization. When starting from an all $+1$ configuration it takes exponentially many time steps to move to the negative phase.*

See Theorem 16 for a formal statement. For graphs which are not such good expanders, the behavior is likely more subtle. Controlling one vertex still is not sufficient to change the phase in polynomial time. However, it suffices to control $o(|V|)$ vertices to change the phase in polynomial time. For example for the $n \times n$ grid we prove the following:

Theorem 7 *Consider the all +1 configuration on the $n \times n$ grid. For $\beta = \Omega(\log n)$ the adversary needs to control linear in n vertices to reach the -1 configuration in a polynomial number of steps. On the other hand, for β a sufficiently large constant, controlling $o(n)$ vertices yields two phases of the stationary measure, and an exponential number of steps is needed to reach the negative phase.*

See Theorems 19, 20, and 23 for the precise results. It is an interesting question to develop a more refined understanding of the interaction between adversarial corruptions and network topology. But our results already set the stage by establishing some important phenomena.

Next we ask the question: Are there macroscopic changes that a dynamic adversary, who does not merely pin all the corrupted vertices, can induce? We study the power of an adversary to prolong polarization. We use the term polarization to refer to the state of having an approximately equal number of +1 and -1 spins, in other words a magnetization close to zero, which is typically an exponentially unlikely configuration. The following theorem gives a tight bound on the size of the corrupted set necessary for the adversary to maintain polarization for longer than polynomial time.

Theorem 8 *For $\beta = \frac{c}{n}$ in the mean field setting, i.e. $G = K_n$, starting from an i.i.d. random configuration, there is a strategy for an adversary controlling $n^{1/2+\epsilon}$ vertices to maintain polarization for a super-polynomial number of time steps (in other words, it prevents convergence to one of the phases in polynomial time).*

To prove this theorem, a more sophisticated and coordinated attack is needed. Our attack is based on assigning the corrupted vertices to the minority spin, until that spin nearly becomes an overwhelming majority. This is formalized in Theorem 24. Note that the random initialization in this result is essential, as corrupting $o(n)$ vertices starting from a global +1 configuration has no desired effect. Understanding the scope and extent of an adversary’s power to shape opinion dynamics is obviously a challenging problem. Our results show that, depending on the goal, an adversary’s strategy can be subtle and hard to detect.

Finally, we address the question: Given a budget of ϵn nodes, which ones should the adversary choose to minimize the magnetization when measured at a specific time in the future? We note that for the ferromagnetic Ising model if we are interested in the long term-behavior (stationary) behavior of the system then it is known that the influence maximization problem is sub-modular, see Bresler et al. (2019). Our main result in this direction is that the sub-modularity extends to the dynamic setting under a sufficiently strong external field.

Theorem 9 *For the corrupted Glauber dynamics on a graph G with external field at least the maximum degree, the influence maximization problem at any time step t is submodular.*

This is formalized in Theorem 33. The proof captures a general phenomenon that submodularity of the local update rule implies global submodularity of the maximization problem.

1.3. Follow-up Work

Since the initial release of our work, several papers have studied our model of adversarial Glauber dynamics. Gaitonde and Mossel (2023) gave a new analysis of a simple algorithm for learning the underlying structure and parameters of Ising models based on independent samples from the

measure. Among their new results includes novel guarantees for learning based on samples from our adversarial Glauber dynamics. [Chen and Mossel \(2024\)](#) uncovered a sharp computational phase transition for the influence maximization problem that we consider. Their methods are based on spectral independence and spatial mixing properties.

1.4. Future Questions

Our framework leaves many interesting questions and research directions.

Is it possible to generalize our result about preventing convergence outside of the mean field setting? It is natural to believe that the conclusion should hold true as long as the graph is sufficiently well-connected, which leads to the following conjecture.

Conjecture 10 *Let G be an expander graph. Then there is a choice of corrupted set that maintains polarization for a super-polynomial number of time steps.*

We expect many interesting phenomena to occur for models that allow anti-ferromagnetic interactions, including spin glasses, where we believe that finding the optimal set (even approximately) is computationally hard. Indeed given that sampling from such models is computationally hard [Sly \(2010\)](#); [Sly and Sun \(2012\)](#), it is natural to expect that influence maximization is also hard.

Another dimension of complexity may be introduced by considering game theoretic models where a number of different players with different objectives may choose different sets. For example, what are the optimal strategies for 3 players in an anti-ferromagnetic Potts models if one wants to maximize the number of green nodes, another the number of blue nodes and the third the number of red nodes at a specific time? We can assume for example that players each have a budget of ϵn nodes and a node is allocated uniformly at random to one of the players that bid for it.

It is interesting to answer analogous questions here, where the study of censored rather than corrupted nodes has also been of recent interest, see e.g. [Moitra et al. \(2021\)](#) and references within.

2. Preliminaries

Here we introduce the remaining main definitions and notation that we will use for the rest of the paper. In order to study the process of how the effect of a corrupted set of vertices permeates through the configuration, we need a process for re-sampling the spins of vertices that agrees with the Ising model distribution. This role will be played by the well-known Glauber dynamics.

Definition 11 (Glauber Dynamics) *The Glauber dynamics is a Markov chain X_t on Ω with the following transition rule. Choose a vertex $v \in V(G)$ uniformly at random. Then $X_{t+1}(u) = X_t(u)$ for every $u \neq v$ and $X_{t+1}(v) = +1$ with probability $\frac{\exp(\beta \sum_{uv \in E(G)} X_t(u))}{\exp(\beta \sum_{uv \in E(G)} X_t(u)) + \exp(-\beta \sum_{uv \in E(G)} X_t(u))}$.*

The transition probability is the conditional probability of the spin at v being $+1$ under the measure μ conditioned on the remaining spins in the graph. It is easy to check that the Glauber dynamics is reversible with respect to the measure μ , and thus has our desired stationary distribution.

The Glauber dynamics is one of the most extensively studied ways to sample from the Ising model distribution. Given the computational hardness of exactly sampling from the measure (see e.g. [Istrail \(2000\)](#)), such a Markov chain is the most common way to sample approximately. There is a rich line of research aimed at analyzing the Glauber dynamics. Through many innovations and

breakthroughs, it is known that in many cases, the chain exhibits a phase transition in the time it takes to reach equilibrium. For a comprehensive overview of techniques see [Martinelli \(1999\)](#) and [Levin and Peres \(2017\)](#).

At a technical level, some of our results will be established using bottlenecks for the magnetization. More generally, a bottleneck is a constraining region of the configuration space Ω that limits the transitions of the Glauber dynamics. The strength of a bottleneck is quantified by the bottleneck ratio.

Definition 12 (Bottleneck Ratio) *Let P be the transition matrix of the Glauber dynamics on Ω , which has stationary distribution μ . Define*

$$Q(\sigma_1, \sigma_2) := \mu(\sigma_1)P(\sigma_1, \sigma_2) \text{ and } Q(A, B) = \sum_{\sigma_1 \in A, \sigma_2 \in B} Q(\sigma_1, \sigma_2).$$

The bottleneck ratio of a subset $S \subset \Omega$ is defined as

$$\Phi(S) := \frac{Q(S, S^c)}{\mu(S)}$$

and the bottleneck ratio of the dynamics is

$$\Phi^* := \min_{\mu(S) \leq \frac{1}{2}} \Phi(S).$$

As alluded above, the bottleneck ratio controls the time it takes for the Glauber dynamics to move from one region of Ω to another, and will be used to bound the number of steps needed for the impact of the corrupted set A to take effect. See [Levin and Peres \(2017\)](#) Theorem 7.4 for a useful result in this spirit.

Notation

For the rest of the paper, we will use the convention that variables with a tilde will refer to the corrupted dynamics (e.g. $\tilde{X}_t, \tilde{\mu}, \tilde{\Phi}$, etc) and variables with no tilde will refer to the corresponding values in the standard dynamics.

3. No Corruption Effect

We first prove “negative” results (in the sense that the corrupted vertices do not change much) in both the high and low temperature regimes. These results will all take place on regular graphs. The first result is in the case of a high temperature (small β) system, where interactions between vertices are close to independent. In this case it is natural to believe that a group of few vertices cannot strongly influence the remaining vertices. This is captured in a quantitative sense by comparing the trajectories between X_t and \tilde{X}_t .

Theorem 13 *Fix $\epsilon > 0$. Suppose $\beta \leq \frac{1}{10d}$, and let G be a d -regular graph. Let X_t be the standard Glauber dynamics, and \tilde{X}_t be ϵn -corrupted Glauber dynamics coupled in the natural way. Let $D_t = \{v \in V(G) \setminus A : X_t(v) \neq \tilde{X}_t(v)\}$ where A is the set of vertices controlled by the adversary. Then $|D_t| \leq 2\epsilon n$ for all times $t \leq e^{\epsilon^2 \beta n}$ with probability $1 - 8\beta\epsilon e^{-\epsilon^2 \beta n}$. In particular, this holds for $t = n^k$ for any $k \in \mathbb{N}$.*

The proof proceeds by comparing the trajectories of the standard dynamics with the corrupted one. We show that when the number of differences is large, there is a bias toward returning to 0. Thus, at any fixed time, it will be unlikely for there to be a large number of differences between the two processes. The first claim is captured by the following lemma, whose proof is in Appendix A.

Lemma 14 *For any $\delta > 0$ let $L_1 = n \left(\delta + \frac{(\epsilon + \delta)d\beta}{1 - d\beta} \right)$. Then for all $L \geq L_1$,*

$$\mathbb{E} [|D_{t+1}| | |D_t| = L] \leq L - \delta.$$

The second claim is the following estimate on the hitting probabilities of a biased random walk. The proof is in Appendix B.

Lemma 15 *Fix $0 < \delta < \frac{1}{4}$. Let W_t be a random walk such that $W_0 = 1$, $\mathbb{P}(W_{t+1} = W_t + 1) = \frac{1}{2} - \delta$, and $\mathbb{P}(W_{t+1} = W_t - 1) = \frac{1}{2} + \delta$. Let $\tau_a = \inf\{t \geq 0 : W_t = a\}$. Then*

$$\mathbb{P}(\tau_a < \tau_0) < 8\delta e^{-2\delta a}.$$

Proof sketch of Theorem 13 Applying Lemma 14 implies that the number of differences between the corrupted and standard dynamics is dominated by a biased random walk whenever it becomes large. Lemma 15 then shows that the difference stays large with very small probability. See Appendix C for the details. ■

We now consider the opposite case of a low temperature (large β) system. In this case, there are strong interactions between adjacent vertices, so it is possible to hope for a large impact on the dynamics while only controlling a small number of vertices. However, as long as the graph is sufficiently well-connected (captured by the notion of edge expansion), then a small constant fraction of the vertices is still insufficient to flip the phase of the configuration efficiently. This is captured by the following theorem.

Theorem 16 *Let G be a d -regular, α -edge-expander graph. There exists a $C(\frac{\alpha}{d})$ such that for $\beta > \frac{1}{2d}C(\frac{\alpha}{d})$ and $\epsilon < \frac{\alpha}{4d}$, the mixing time of the ϵn -corrupted Glauber dynamics is at least $C \exp(\frac{1}{4}\alpha\beta n)$. In particular, if $\alpha = \Theta(d)$ then this holds for $\beta = \Omega(\frac{1}{d})$.*

Remark 17 *In the above theorem, the corrupted set is fixed to $+$, so the mixing time refers to the mixing time of standard Glauber dynamics with boundary conditions.*

Proof sketch The proof is a generalization of the standard bottleneck calculation for the Ising model on the complete graph. See Appendix D for the details. ■

4. Forcing on the Grid

In this section we give a series of “positive” results in which the corruption has a macroscopic effect the dynamics in polynomial time. These results will take place on the $n \times n$ square lattice. As opposed to the prior results which are for general graphs satisfying certain nice properties, the analysis of Glauber dynamics on the grid rely on the specific geometric structure of the graph. The common question we study is: starting from the $+1$ configuration, how many nodes need to be corrupted in order to force the dynamics to the -1 configuration in a polynomial (in n) number of steps?

Remark 18 *By the monotonicity of the measure μ , we only need to consider the case in which the corrupted vertices are fixed to the -1 spin.*

To start, we consider the case of zero temperature, which is closely tied to the study of bootstrap percolation (see [Balogh and Pete \(1998\)](#) for further background on this related problem).

Theorem 19 *Let $\beta = \infty$ and G be the square lattice with vertex set $[n]^2$. Suppose the configuration is initialized to be identically $+1$. Then:*

1. *The m -corrupted Glauber dynamics can never reach the -1 configuration for $m < n$.*
2. *Let $A = \{(i, i) : i \in [n]\}$. Then the A -corrupted Glauber dynamics will reach the -1 configuration in expected time $O(n^5)$.*

Proof sketch Statement 1 follows from a well-known perimeter argument in bootstrap percolation. We prove statement 2 by noting that once the triangle $T_k = \{(i, j) : i < j, j < k\}$ has all -1 s, the configuration will never change. The time for T_k to become -1 is estimated by comparison to a simple random walk. See [Appendix E](#) for the details. ■

As a corollary of the above theorem, we can study the case of slowly decreasing temperature.

Theorem 20 *Let $\beta \geq 4 \log n$ and G be the square lattice with vertex set $[n]^2$. Suppose the initialize with the identically $+1$ configuration. Then for $A = \{(i, i) : i \in [n]\}$, the A -corrupted Glauber dynamics will reach the -1 configuration in expected time $O(n^5)$.*

Proof sketch The main observation is that at sufficiently low temperature, the behavior of the chain is unlikely to differ from that of a zero temperature chain during a polynomial number of steps. The details are contained in [Appendix F](#). ■

Remark 21 *It follows from the proof that once the all -1 state is reached, the configuration remains in this state for at least n^8 time in expectation. In particular, since the time taken to reach the state is $O(n^5)$ on average, the configuration will remain in the -1 state for $1 - o(1)$ fraction of time.*

Remark 22 *It does not follow from the perimeter argument cited previously that $n - 1$ vertices does not suffice for a polynomial time phase flip. Indeed, given $n - 1$ vertices as the set A , a deviation at the remaining diagonal site will occur in expected polynomial time, after which the global spin configuration will flip in polynomial time once again.*

Finally, to handle the case of low (but constant) temperature, considerably more work is needed. In this case, the fluctuations of the model prevent us from converging to the -1 configuration. However, the configuration converges to the -1 phase, which is the distribution on configurations under -1 boundary conditions. In this setting we show a transition in behavior when the size of the corrupted set moves from sub-linear to linear.

Theorem 23 *Let G be the $n \times n$ grid, and let β be a sufficiently large constant. Starting from the all $+1$ configuration, the An -corrupted Glauber dynamics vertices converges to the -1 phase in sub-exponential time. Conversely, the $o(n)$ -corrupted Glauber dynamics mixes in exponential time.*

Proof sketch The analysis in this case builds on a series of results by Martinelli which use the geometric structure of the square grid to analyze the standard Glauber dynamics under various boundary conditions. We use these results to show that fixing a sub-linear number of vertices does not significantly change the stationary measure. The proof is in Appendix G. \blacksquare

5. Efficient Polarization

We next present a new question that demonstrates the flexibility of our new model. As opposed to the above results in which the goal was to influence the overall state of the network as much as possible in a specified direction, we now consider the goal of introducing as much disorder as possible. If the spins of the configuration are initialized randomly, there will be a strong misalignment within the system. We desire the effect of oscillating the configuration repeatedly, perpetuating the polarization of the configuration. In this case, it is also natural to ask how much of the network must be corrupted to induce this behavior.

Theorem 24 *Let $G = K_n$ with an i.i.d. initial configuration in $n^{1/2+\epsilon}$ -corrupted Glauber dynamics. With probability at least $(1 - e^{-\beta n^{1+\epsilon}})^T - e^{-n^{2\epsilon}/32}$, the adversary can force $m(\tilde{X}_t)$ to enter $[-2\sqrt{n}, 2\sqrt{n}]$ and fluctuate within this interval up to time T before $|m(\tilde{X}_t)|$ ever reaches $n^{1/2+\epsilon}$. In particular, for $\beta = \frac{c}{n}$ the dynamics will not converge in polynomial time.*

Remark 25 *Note that a randomly initialized configuration has fluctuations in magnetism of order \sqrt{n} , so to influence the network the control of the adversary must dominate this fluctuation.*

Proof At initialization, set the spins on A to be identically $+1$. Let $M_t = m(\tilde{X}_t(V(G) \setminus A))$. Whenever $M_t \geq -\frac{1}{2}n^{1/2+\epsilon}$ we have the following estimates. Suppose that a vertex v is selected by the dynamics and $\sigma(v) = -1$. Then the conditional probability that the spin is updated to $+1$ is (recall Definition 11)

$$\frac{e^{\beta(|A|+M_t+1)}}{e^{\beta(|A|+M_t+1)} + e^{-\beta(|A|+M_t+1)}} \geq \frac{\exp(\frac{1}{2}\beta n^{1/2+\epsilon})}{\exp(\frac{1}{2}\beta n^{1/2+\epsilon}) + \exp(-\frac{1}{2}\beta n^{1/2+\epsilon})} =: p.$$

Similarly, if v is selected with $\sigma(v) = +1$, the conditional probability that the spin is updated to -1 is

$$\frac{e^{-\beta(|A|+M_t-1)}}{e^{\beta(|A|+M_t-1)} + e^{-\beta(|A|+M_t-1)}} \leq \frac{\exp(-\frac{1}{2}\beta n^{1/2+\epsilon})}{\exp(\frac{1}{2}\beta n^{1/2+\epsilon}) + \exp(-\frac{1}{2}\beta n^{1/2+\epsilon})} =: q.$$

Thus, in this regime M_t stochastically dominates a lazy random walk with $\frac{p}{2}$ probability of a $+1$ step and $\frac{q}{2}$ probability of a -1 step. Thus, started from a configuration with magnetization $\geq -\frac{1}{4}n^{1/2+\epsilon}$, a standard calculation similar to Lemma 15 shows that the probability that M_t hits $n^{1/2}$ before $-\frac{1}{2}n^{1/2+\epsilon}$ is at least

$$\frac{\left(\frac{q}{p}\right)^{-\frac{1}{4}n^{1/2+\epsilon}} - \left(\frac{q}{p}\right)^{-\frac{1}{2}n^{1/2+\epsilon}}}{\left(\frac{q}{p}\right)n^{1/2} - \left(\frac{q}{p}\right)^{-\frac{1}{2}n^{1/2+\epsilon}}} \geq 1 - e^{-\frac{1}{4}\beta n^{1+2\epsilon}}.$$

In particular, the probability of failure is exponentially small in $\beta n^{1+2\epsilon}$. The same calculation shows that when started from a configuration with magnetization $\geq -n^{1/2}$, the probability that M_t hits $n^{1/2}$ before $-2n^{1/2}$ is at least $1 - \exp(-\beta n^{1+\epsilon})$.

By symmetry, the analogous statements hold for $M_t \leq \frac{1}{2}n^{1/2+\epsilon}$ when the spins on A are set to -1 . In particular, started from a configuration with magnetization $\leq \frac{1}{4}n^{1/2+\epsilon}$, the probability that M_t hits $-n^{1/2}$ before $\frac{1}{2}n^{1/2+\epsilon}$ is at least $1 - e^{-\frac{1}{4}\beta n^{1+2\epsilon}}$. Also, started from a configuration with magnetization $\leq n^{1/2}$, the probability that M_t hits $-n^{1/2}$ before $2n^{1/2}$ is at least $1 - \exp(-\beta n^{1+\epsilon})$. Thus, by flipping the corrupted set whenever the magnetization reaches $-n^{1/2}$ from $n^{1/2}$ and vice versa, the number of iterations that the magnetization of the system oscillates between $-2\sqrt{n}$ and $2\sqrt{n}$ under the corruption of A dominates a geometric random variable with mean $e^{\beta n^{1+\epsilon}}$.

To conclude it suffices to show that M_0 is between $-\frac{1}{4}n^{1/2+\epsilon}$ and $\frac{1}{4}n^{1/2+\epsilon}$ with high probability, so the adversary has control at initialization. Since each spin is independent and uniform, $M_0 \stackrel{d}{=} 2\text{Bin}(n - |A|, \frac{1}{2}) - (n - |A|)$. An application of Hoeffding's Inequality implies

$$\mathbb{P}\left(|M_0| > \frac{1}{4}n^{1/2+\epsilon}\right) = \mathbb{P}\left(|2\text{Bin}(n - |A|, \frac{1}{2}) - (n - |A|)| > \frac{1}{4}n^{1/2+\epsilon}\right) \leq e^{-n^{2\epsilon}/32}.$$

■

6. Maximizing the Adversary's Effect

To conclude, we present a series of results about the optimization problem from the perspective of the adversary. Namely, given a budget of m nodes, we study the optimal strategy for the adversary to implement to induce the maximum effect at time t . We start with a result about the behavior at high temperature. In the remainder of this section, for v a vertex in a graph G , let $N(v)$ denote the neighbors of v , $d(v) = |N(v)|$ denote the degree, and $d_S(v)$ denote the number of edges from v to the set S .

Proposition 26 *Let G be a graph of maximum degree d and $M(\epsilon)$ be the ϵ -mixing time for the Glauber dynamics for the Ising model on G with inverse temperature β . Let \tilde{X}_t^A be an instance of the A -corrupted Glauber dynamics for $|A| = m$. For $\beta \leq \frac{1}{4d^2m}$, $\epsilon = \frac{\beta}{4n}$ and $t \geq M(\epsilon)$, $\mathbb{E}\left[\left|\{v : \tilde{X}_t^A(v) = +1\}\right|\right]$ is maximized when A are the m vertices with largest degree to the uncorrupted vertices, $V \setminus A$.*

Proof sketch The heuristic in this high temperature regime is that the Glauber dynamics updates are well approximated by linear functions, and thus the degree of the corrupted vertices is the dominant contributing factor. The proof in Appendix H formalizes this heuristic. ■

Remark 27 *For graphs G and temperatures β for which the Glauber dynamics mixes rapidly in polynomial time, we have that $M(\epsilon)$ for this choice of epsilon is also polynomial in n . In particular, the best choice for the adversary is to choose the highest degree vertices, and this choice manifests itself quickly.*

In situations where the explicit optimizer is not as easy to compute, submodularity is a desirable property for the maximization of set functions. It is analogous to convexity, and thus allows for efficient algorithms for approximate optimization, see for example the work of Feige et al. (2011). It was shown by Bresler et al. (2019) using the GHS inequality that for the ferromagnetic Ising

model, the influence maximization problem for the equilibrium measure is submodular. It is natural to ask whether or not this property extends to the dynamic, polynomial-time setting as well. The following observations and results begin to answer this question.

Definition 28 (Submodularity) *We say that a function $f : V \rightarrow \mathbb{R}$ is submodular if for all $S, T \subseteq V$, $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$.*

Observation *The function $f_v : 2^{N(v)} \rightarrow [0, 1]$ defined by $f_v(S) = \frac{e^{2|S|-d(v)}}{e^{2|S|-d(v)} + e^{-2|S|+d(v)}}$ is not submodular.*

The function f_v is the local update of the Glauber dynamics at vertex v when the set of neighbors of v with $+1$ spin is exactly S . This function fails to be submodular when the sets S and T are too small. Thus, it immediately follows by considering a 10-regular graph and disjoint S and T of size 2 that the function $\mathbb{E} \left[\left\{ v : \tilde{X}_t^S(v) = +1 \right\} \right]$ cannot be a submodular function of S in general for all graph G and all times t . However, we know that as $t \rightarrow \infty$ the expectation approaches a submodular function. Thus, it is natural to ask whether or not there is a characterization for which graphs G , or for which times t , that the property of submodularity holds. We provide answers in two situations.

6.1. Approximate Submodularity

Definition 29 (Approximate Submodularity) *We say that a function $f : V \rightarrow \mathbb{R}$ is ϵ approximately submodular if for all $S, T \subseteq V$, $f(S) + f(T) \geq f(S \cup T) + f(S \cap T) - \epsilon$.*

The following proposition gives a sufficient condition for the optimization problem to be approximately submodular. This gives hope for approximately solving the maximization.

Proposition 30 *Let $M(\epsilon)$ be the ϵ -mixing time of the Glauber dynamics on a graph G with parameter β . In particular, for $t \geq M(\epsilon)$ we have $d_{\text{TV}}(\tilde{\mu}_t^S, \mu) < \epsilon$ where $\tilde{\mu}_t^S$ is the distribution of \tilde{X}_t^S and μ is the Ising model Gibbs measure. Then for $t \geq M(\epsilon)$, $\mathbb{E} \left[\left\{ v : \tilde{X}_t^S(v) = +1 \right\} \right]$ is $4\epsilon n$ -approximately submodular as a function of S .*

Remark 31 *In the above proposition the adversary's optimal strategy is to assign all corrupted vertices to $+$, so the mixing time refers to the standard mixing time with boundary conditions.*

Proof Note that the quantity $\left| \left\{ v : \tilde{X}_t^S(v) = +1 \right\} \right|$ is a random variable bounded below by 0 and above by n . The event where X_t^S differs from μ has measure at most ϵ , and the random variable differs by at most n on this event, so the expectation under the two measures differs by at most ϵn . Applying this discrepancy to each term in the definition of submodularity yields the desired inequality. ■

Note that the above proposition in particular implies for rapidly mixing chains, the maximization problem becomes approximately submodular within a polynomial in n number of time steps.

6.2. Submodularity Under Strong External Field

We conclude by showing that in the presence of a strong enough external field, the maximization problem is in fact submodular for all time steps. The following lemma establishes that the local update becomes submodular with a strong external field. See Appendix I for the proof.

Lemma 32 *Let h be at least the maximum degree of G . Let $f_v : 2^{N(v)} \rightarrow [0, 1]$ be defined by $S \mapsto \frac{\exp(\beta(2|S|-d_v+h))}{\exp(\beta(2|S|-d_v+h))+\exp(\beta(-2|S|+d_v-h))}$. Then f_v is submodular.*

Theorem 33 *For any $S \subseteq V$, let \tilde{X}_t^S be the S -corrupted Glauber dynamics with external field h at least the maximum degree of G . The function $\mathbb{E} \left[\left| \{v : \tilde{X}_t^S(v) = +1\} \right| \right]$ is submodular function of S for every t .*

Proof sketch The proof is an extension of local submodularity (given by Lemma 32) to global submodularity, using a coupling closely related to one introduced by Mossel and Roch (2010). See J for the details of the coupling. ■

Remark 34 *The above proof in general shows that for any submodular local update f_v , running the Glauber dynamics will result in a submodular optimization problem for all times t .*

Acknowledgments

BC was supported by the NSF Graduate Research Fellowship Program. EM and CS were partially supported by Vannevar Bush Faculty Fellowship ONR-N00014-20-1-2826 and by Simons Investigator award (622132) EM was also partially supported by NSF award CCF 1918421. AM and CS were also partially supported by AM's ONR YIP grant. AM was also partially supported by a David and Lucile Packard Fellowship.

References

- Kristine Eia S Antonio, Chryslie Margus N Pinol, and Ronald S Banzon. An ising model approach to malware epidemiology. *arXiv preprint arXiv:1007.4938*, 2010.
- Simone Baldassarri, Anna Gallo, Vanessa Jacquier, and Alessandro Zocca. Ising model on clustered networks: A model for opinion dynamics. *arXiv preprint arXiv:2212.12536*, 2022.
- József Balogh and Gábor Pete. Random disease on the square grid. *Random Structures & Algorithms*, 13(3-4):409–422, 1998.
- Guy Bresler, Frederic Koehler, and Ankur Moitra. Learning restricted boltzmann machines via influence maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 828–839, 2019.
- Zongchen Chen and Elchanan Mossel. Influence maximization in ising models, 2024.
- Martin Dyer, Alistair Sinclair, Eric Vigoda, and Dror Weitz. Mixing in time and space for lattice spin systems: a combinatorial view. *Random Structures Algorithms*, 24(4):461–479, 2004.

- Uriel Feige, Vahab S Mirrokni, and Jan Vondrák. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- Jason Gaitonde and Elchanan Mossel. A unified approach to learning ising models: Beyond independence and bounded width, 2023.
- Roy J. Glauber. Time-dependent statistics of the ising model. *Journal of Mathematical Physics*, 4(2):294–307, 1963.
- Sorin Istrail. Statistical mechanics, three-dimensionality and np-completeness: I. universality of intracatability for the partition function of the ising model across non-planar surfaces. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 87–96, 2000.
- Libin Jiang, Mathieu Leconte, Jian Ni, R. Srikant, and Jean Walrand. Fast mixing of parallel glauber dynamics and low-delay csma scheduling. *IEEE Transactions on Information Theory*, 58(10):6541–6555, 2012.
- Stuart A Kauffman. Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of theoretical biology*, 22(3):437–467, 1969.
- D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *Proc. 9th ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, 2003.
- Cristina Gabriela Aguilar Lara, Eduardo Massad, Luis Fernandez Lopez, Marcos Amaku, et al. Analogy between the formulation of ising-glauber model and si epidemiological model. *Journal of Applied Mathematics and Physics*, 7(05):1052, 2019.
- David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- Adam Lipowski. Ising model: Recent developments and exotic applications. *Entropy*, 24(12):1834, 12 2022. ISSN 1099-4300.
- Shihuan Liu, Lei Ying, and Srinivas Shakkottai. Influence maximization in social networks: An ising-model-based approach. In *2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 570–576, 2010.
- Jacek Majewski, H. Li, and Jurg Ott. The ising model in physics and statistical genetics. *American journal of human genetics*, 69 4:853–62, 2001.
- Fabio Martinelli. Lectures on glauber dynamics for discrete spin models. In *Lectures on probability theory and statistics*, pages 93–191. Springer, 1999.
- Fabio Martinelli and Enzo Olivieri. Approach to equilibrium of glauber dynamics in the one phase region: I. the attractive case. *Communications in Mathematical Physics*, 161(3):447–486, 1994.
- Fabio Martinelli and Fabio Lucio Toninelli. On the mixing time of the 2d stochastic ising model with “plus” boundary conditions at low temperature. *Communications in Mathematical Physics*, 296(1):175–213, 2010.

- Ankur Moitra, Elchanan Mossel, and Colin P Sandon. Learning to sample from censored markov random fields. In *Conference on Learning Theory*, pages 3419–3451. PMLR, 2021.
- Andrea Montanari and Amin Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences*, 107(47):20196–20201, 2010.
- E. Mossel and S. Roch. Submodularity of influence in social networks: From local to global. *SIAM J. Comput.*, 39(6):2176–2188, 2010. URL <http://link.aip.org/link/?SMJ/39/2176>.
- Rudolf Peierls. On ising’s model of ferromagnetism. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 32, pages 477–481. Cambridge University Press, 1936.
- M Leticia Rubio Puzzo and Ezequiel V Albano. The damage spreading method in monte carlo simulations: A brief overview and applications to confined magnetic materials. 2008.
- A. Sly. Computational transition at the uniqueness threshold. In *Foundations of Computer Science (FOCS)*, pages 287–296, 2010.
- Allan Sly and Nike Sun. The computational hardness of counting in two-spin models on d -regular graphs. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 361–369. IEEE, 2012.
- H Eugene Stanley, Dietrich Stauffer, Janos Kertesz, and Hans J Herrmann. Dynamics of spreading phenomena in two-dimensional ising models. *Physical review letters*, 59(20):2326, 1987.
- Dror Weitz. Combinatorial criteria for uniqueness of gibbs measures. *Random Struct. Algorithms*, 27:445–475, 2005.

Appendix A. Proof of Lemma 14

Proof Suppose that $|D_t| = L$. In order for D_t to decrease, we need the Glauber dynamics to select a vertex in the difference set. This happens with probability $\frac{L}{n}$. Then to decrease, the updates must match, which occurs with probability $1 - \left| \mu_{X_t}(\sigma_v = +) - \mu_{\tilde{X}_t}(\sigma_v = +) \right|$. In order for D_t to increase, it needs to select a vertex outside the difference set with probability $1 - \epsilon - \frac{L}{n}$, and produce updates that do not match with probability $\left| \mu_{X_t}(\sigma_v = +) - \mu_{\tilde{X}_t}(\sigma_v = +) \right|$. The remaining ϵ probability is when Glauber selects a vertex in the corrupted set, at which D_t does not change.

To improve the dependence on d and β and obtain the desired L_1 , we further distinguish the cases by the number of differences in the neighborhood of the selected vertex v to compute the above probabilities. Let

$$V_t(k) = \{v : N_{X_t}(v) \text{ and } N_{\tilde{X}_t}(v) \text{ contain } k \text{ differences}\}.$$

By double counting,

$$\sum_{k \geq 0} k \cdot |V_t(k)| \leq (\epsilon n + L)d.$$

Now, if the neighborhood of v has k differences, the difference in probabilities is

$$\frac{1}{2} \tanh(\beta(s+k)) - \frac{1}{2} \tanh(\beta(s-k)) \leq \tanh(k\beta).$$

Thus we want

$$\begin{aligned} & \frac{1}{n} \sum_{k \geq 0} |D_t \cap V_t(k)| \cdot (1 - \tanh(k\beta)) - |D_t^c \cap V_t(k)| \cdot \tanh(k\beta) \\ &= \frac{1}{n} \sum_{k \geq 0} |D_t \cap V_t(k)| - |V_t(k)| \tanh(k\beta) > \delta. \end{aligned}$$

Since $\tanh(x) < x$ for $x > 0$, it suffices for

$$\frac{L}{n} - \frac{1}{n} \sum_{k \geq 0} |V_t(k)| \cdot k\beta \geq \frac{L}{n} - \frac{\beta(\epsilon n + L)d}{n} = \frac{L(1-d\beta)}{n} - \epsilon d\beta > \delta.$$

This holds when

$$L \geq \frac{n(\delta + \epsilon d\beta)}{1 - d\beta} = n \left(\delta + \frac{(\epsilon + \delta)d\beta}{1 - d\beta} \right).$$

■

Appendix B. Proof of Lemma 15

Proof By considering the martingale $M_t = \left(\frac{\frac{1}{2} + \delta}{\frac{1}{2} - \delta}\right)^{W_t}$ and the stopping time $\tau = \min\{\tau_a, \tau_0\}$ we can compute

$$\frac{\frac{1}{2} + \delta}{\frac{1}{2} - \delta} = \mathbb{E}[M_0] = \mathbb{E}[M_\tau] = \mathbb{P}(\tau_0 < \tau_a) + \mathbb{P}(\tau_a < \tau_0) \cdot \left(\frac{\frac{1}{2} + \delta}{\frac{1}{2} - \delta}\right)^a$$

Rearranging gives

$$\mathbb{P}(\tau_a < \tau_0) = \frac{\frac{\frac{1}{2} + \delta}{\frac{1}{2} - \delta} - 1}{\left(\frac{\frac{1}{2} + \delta}{\frac{1}{2} - \delta}\right)^a - 1} < \frac{8\delta}{(1+4\delta)^a - 1} < 8\delta e^{-2\delta a}.$$

■

Appendix C. Proof of Theorem 13

Proof Set $\delta = \beta\epsilon$ and let L_1 be as in Lemma 14. Define $T_0 = \inf\{t : |D_t| > L_1\}$ and $S_0 = \inf\{t > T_0 : |D_t| = L_1\}$. Inductively define $T_k = \inf\{t > S_{k-1} : |D_t| > L_1\}$ and $S_k = \inf\{t > T_k : |D_t| = L_1\}$, so that the time intervals $[T_k, S_k]$ capture the excursions of $|D_t|$ above L_1 . For each excursion, associate a new process $(W_t^{(k)})_{t=0}^{S_k - T_k}$ started from 0 that increases by 1 with probability $\frac{1}{2} - \delta$ and decreases by 1 with probability $\frac{1}{2} + \delta$ whenever $|D_t|$ moves. By Lemma 14 and the fact

that $|D_{t+1}| \in [|D_t| - 1, |D_t| + 1]$, we can couple the excursion $(|D_t|)_{t=T_k}^{S_k}$ with $W_t^{(k)}$ such that $|D_t| \leq W_{t-T_k}^{(k)} + L_1$. Since each $W_t^{(k)}$ is an identically distributed biased random walk we know that for each k , by Lemma 15,

$$\mathbb{P}\left(\exists t \leq S_k - T_k : W_t^{(k)} \geq \epsilon n\right) < 8\delta e^{-2\delta\epsilon n}.$$

Now, for any time fixed time t , if $|D_t| \geq L_1 + \epsilon n$, then by the invariant above it must be coupled to a biased walk satisfying $W_{t'}^{(k)} \geq \epsilon n$ for some corresponding $k \leq t$ and t' . Taking a union bound over all $k \leq e^{\epsilon^2\beta n}$,

$$\begin{aligned} \mathbb{P}\left(\exists t \leq e^{\epsilon^2\beta n} : |D_t| \geq 2\epsilon n\right) &< \mathbb{P}\left(\exists t \leq e^{\epsilon^2\beta n} : |D_t| \geq L_1 + \epsilon n\right) \\ &\leq \mathbb{P}\left(\exists k \leq e^{\epsilon^2\beta n}, t' : W_{t'}^{(k)} \geq \epsilon n\right) < 8\epsilon\beta e^{-\epsilon^2\beta n} \end{aligned}$$

which is the desired inequality. ■

Appendix D. Proof of Theorem 16

Proof The proof proceeds by exhibiting a bottleneck in the dynamics of the process. Suppose G a d -regular expander with $d \geq \min_{0 \leq |S| \leq \frac{n}{2}} \frac{|\partial S|}{|S|} \geq \alpha$. Define $A_k = \{\sigma : |\{v : \sigma_v = +\}| = k\}$. For $\sigma \in A_k$, we know $dk \geq |\partial S| \geq \alpha k$ where $S = \{v : \sigma_v = +\}$. Note that

$$\sum_{u \sim v} \sigma_u \sigma_v = (|E(G)| - |\partial S|) - |\partial S| \in \left[\frac{nd}{2} - 2dk, \frac{nd}{2} - 2\alpha k\right].$$

Now, consider a dynamics that is corrupted with ϵn vertices fixed to $+$. Here, with only $n(1 - \epsilon)$ free vertices, we adjust the bottleneck to $k^* = \frac{n(1-2\epsilon)}{2}$. Note that the number of configurations with $k < k^*$ is less than that of $k > k^*$, and also the configurations with $k < k^*$ have lower probabilities due to the fixed set of $+$ vertices. Thus, $\tilde{\mu}\left(\bigcup_{k < k^*} A_k\right) < \frac{1}{2}$. With $k^* + \epsilon n$ vertices labeled $+'$ s,

$$\begin{aligned} \tilde{\mu}(A_{k^*}) &\leq \frac{1}{Z_c} \cdot \binom{n(1-\epsilon)}{\frac{n(1-2\epsilon)}{2}} \exp\left(\beta n \left(\frac{d}{2} - \alpha\right)\right) =: \frac{1}{Z_c} \cdot \tilde{a}_{k^*} \\ \tilde{\mu}\left(\bigcup_{k < k^*} A_k\right) &\geq \max_{k > k^*} \frac{1}{Z_c} \cdot \binom{n(1-\epsilon)}{k} \exp\left(\beta \left(\frac{nd}{2} - 2d(k + \epsilon n)\right)\right) =: \max_{k < k^*} \frac{1}{Z_c} \cdot \tilde{a}_k \end{aligned}$$

By the same asymptotic approximations, and writing $k = cn(1 - \epsilon)$,

$$\begin{aligned} \log \tilde{a}_{cn(1-\epsilon)} &= \beta nd \left(\frac{1}{2} - 2(c(1 - \epsilon) + \epsilon)\right) - n(1 - \epsilon)[c \log c + (1 - c) \log(1 - c) + o(1)] \\ \log \tilde{a}_{k^*} &\leq n(1 - \epsilon) \log 2 + \beta n \left(\frac{d}{2} - \alpha\right) \end{aligned}$$

The first expression has the same critical point at $c^* = \frac{1}{1+e^{2\beta d}}$, so we can bound the log of the bottleneck ratio as

$$\begin{aligned} \log \tilde{\Phi}^* &\leq n \left[\beta \left(\frac{d}{2} - \alpha \right) + (1 - \epsilon) \log 2 \right] \\ &\quad - n \left[\beta d \left(\frac{1}{2} - 2(c^*(1 - \epsilon) + \epsilon) \right) - (1 - \epsilon) [c^* \log c^* + (1 - c^*) \log(1 - c^*)] \right] \\ &\leq n[(1 - \epsilon)(\log 2 + c^* \log c^* + (1 - c^*) \log(1 - c^*)) + \beta(-\alpha + 2d(c^*(1 - \epsilon) + \epsilon))]. \end{aligned}$$

The rightmost term is at most $-\frac{1}{2}n\alpha\beta$ as long as $c^* < \frac{\alpha}{4d(1-\epsilon)}$ and $\epsilon < \frac{\alpha}{4d}$. The first condition translates to $\beta > \frac{1}{2d} \log(\frac{8d}{\alpha} - 1)$.

To handle the remaining term, we set $C := 2d\beta$. We want to ensure that

$$\log 2 + c^* \log c^* + (1 - c^*) \log(1 - c^*) \leq \frac{\alpha\beta}{4} = \frac{C\alpha}{8d}.$$

Taking the derivative of the left hand side with respect to C , we get the expression $\frac{Ce^C}{(1+e^C)^2}$. Notice that as $C \rightarrow \infty$, this derivative tends to zero, while $\frac{d}{dC} \frac{C\alpha}{8d} = \frac{\alpha}{8d}$. Thus, there exists some $C(\frac{\alpha}{d})$ such that for all $C > C(\frac{\alpha}{d})$, the desired inequality holds. In particular, for $\beta \geq \frac{C(\alpha/d)}{2d}$ we have $\log \tilde{\Phi}^* \leq -\frac{\alpha\beta}{4}$.

This gives the desired bottleneck in the corrupted case. The result for mixing time follows from a standard inequality (i.e. [Levin and Peres \(2017\)](#) Theorem 7.4). \blacksquare

Appendix E. Proof of Theorem 19

Proof Note that at zero temperature, the vertices that can eventually transition to a -1 state is exactly the bootstrap percolation set of the corrupted vertices. Thus, the first claim follows from a perimeter argument ([Balogh and Pete \(1998\)](#) Folklore Fact).

For the second claim, we relate the progress of the configuration toward the -1 configuration to a random walk. Let $C_j = \{(i, j) : i < j\}$ and $T_j = \bigcup_{k \leq j} C_k$. Notice that once the configuration is identically -1 on T_j , then this portion of the configuration is frozen in this state. Define $S_j = \inf\{t \geq 0 : T_j \text{ is identically } -1 \text{ at time } t\}$. We bound $S_{j+1} - S_j$ by comparison to a random walk. Notice that C_{j+1} must flip to the $-$ state in the order $(j, j+1), (j-1, j+1), \dots, (1, j+1)$. Moreover, if $(j, j+1), \dots, (k, j+1)$ are all in the -1 state, then $(j, j+1), \dots, (k+1, j+1)$ are all frozen in the -1 state. Thus, the quantity $X_t = \inf\{k : (k, j+1) \text{ is } - \text{ in } \sigma_t\}$ is a simple random walk on $[j]$ reflected at j and absorbed at 1. By the Glauber dynamic update, $\mathbb{P}(X_{t+1} = X_t + 1) = \mathbb{P}(X_{t+1} = X_t - 1) = \frac{1}{2n^2}$. Starting from j the expected time to absorption of X_t is j^2 , so $\mathbb{E}[S_{j+1} - S_j] = 2n^2j^2$. Thus,

$$\mathbb{E}[S_n] = \sum_{j=1}^{n-1} \mathbb{E}[S_{j+1} - S_j] = \sum_{j=1}^{n-1} 2n^2j^2 = O(n^5).$$

By symmetry, the same argument applies to the lower triangle below the diagonal, and the hitting time of the -1 configuration is the maximum of the upper triangle and lower triangle, yielding the desired bound. \blacksquare

Appendix F. Proof of Theorem 20

Proof The conditional probability of a vertex whose neighborhood has sum $s > 0$ to be updated to a $-$ state is $\frac{e^{-\beta s}}{e^{\beta s} + e^{-\beta s}} = \frac{n^{-s}}{n^s + n^{-s}} = \frac{1}{n^{8s+1}} < \frac{1}{n^8}$. By symmetry the conditional probability of a vertex whose neighborhood has sum $s < 0$ to be updated to a $+1$ state is also $< \frac{1}{n^8}$. Denote these events as “deviations” from the $\beta = \infty$ case studied above in Proposition 19. Let T be the first time that the configuration is identically -1 in the setting of Proposition 19, and let N be the time of the first deviation in the process. Then N dominates N' , a $\text{Geom}(n^{-8})$ random variable that is independent of T . We have the estimate

$$\mathbb{P}(N < n^6) \leq \mathbb{P}(N' < n^6) = 1 - \left(1 - \frac{1}{n^8}\right)^{n^6} < n^{-2}.$$

We can bound

$$\begin{aligned} \mathbb{P}(N' < T) &= \sum_{m < n^6} \mathbb{P}(m < T) \mathbb{P}(N' = m) + \sum_{m \geq n^6} \mathbb{P}(m < T) \mathbb{P}(N' = m) \\ &\leq \mathbb{P}(N < n^6) + \sum_{m \geq n^6} \mathbb{P}(n^6 < T) \mathbb{P}(N' = m) \\ &\leq n^{-2} + \frac{\mathbb{E}[T]}{n^6} \cdot \mathbb{P}(N' \geq n^6) \\ &\leq n^{-2} + Cn^{-1} \end{aligned}$$

where in the last step we used Markov’s Inequality and Proposition 19. Thus, with probability $1 - o(1)$, we do not deviate before reaching the -1 configuration. In the event that a deviation occurs, we “reset”, and by monotonicity the expected time to reach the -1 configuration from this new configuration must also be bounded by $O(n^5)$. Conditioning on the number of runs until we reach the -1 state without a deviation, we obtain the claimed result. \blacksquare

Appendix G. Proof of Theorem 23

Proof As shown in [Martinelli and Toninelli \(2010\)](#) imposing -1 boundary conditions on the $n \times n$ grid yields a mixing time of at most $\exp(cn^\epsilon)$. Thus, the all $+1$ configuration can be forced into the -1 phase in sub-exponential time with control of $O(n)$ vertices.

Conversely, with control of $o(n)$ vertices, the bottleneck that yields exponential time mixing for the square grid remains, so it is not possible to force the $+1$ phase to the -1 phase in sub-exponential time. We will make repeated use of the following theorem from [Martinelli \(1999\)](#), Proposition 6.3.

Theorem For β sufficiently large, and $\Delta \in (-m^*(\beta), m^*(\beta))$,

$$\lim_{n \rightarrow \infty} -\frac{1}{\beta(2n+1)} \log(\mu(\{\sigma : m(\sigma) = \Delta n^2\})) = c_\Delta > 0$$

where $m^*(\beta)$ is a constant known as the spontaneous magnetization.

We use this to construct a bottleneck in the corrupted dynamics. Note that by symmetry, in the standard dynamics with $S = \{\sigma : m(\sigma) > 0\}$, we have $\mu(S) = \frac{1 - e^{-c\beta(2n+1)}}{2} < \frac{1}{2}$ and $\Phi(S) < 4e^{-c\beta(2n+1)}$ so the bottleneck ratio is exponentially small in n . In the corrupted dynamics, we once again have $\tilde{\mu}(S) < \frac{1}{2}$.

The idea of the argument is that since the maximum degree of G is 4, the corrupted set of vertices can affect at most $o(n)$ edges, which means the probability of each state σ can change by at most a factor of $e^{o(n)}$. Thus, we have that $\tilde{\Phi}(S) < 4e^{-c\beta(2n+1)+o(n)}$, which also yields an exponential mixing time of at least $e^{c\beta(2n+1)-o(n)}$. This is formalized below to conclude the proof.

We begin with a lemma that controls the change in the normalization constant of the two distributions.

Lemma 35 $Z(\beta)e^{-o(n)} \leq \tilde{Z}(\beta) \leq Z(\beta)$.

Proof The upper bound is clear, since the configurations in the corrupted model are a strict subset of the configurations for the standard dynamics, so we prove the lower bound. For $\tau \in \{+1, -1\}^A$, let $\Omega_\tau = \{\sigma : \sigma(A) = \tau\}$. We can re-write

$$Z(\beta) = \sum_{\Omega} e^{-\beta H(\sigma)} = \sum_{\tau} \sum_{\Omega_\tau} e^{-\beta H(\sigma)} \leq \sum_{\tau} e^{4\beta|A|} \tilde{Z}(\beta) \leq 2^{|A|} e^{4\beta|A|} \tilde{Z}(\beta) = e^{o(n)} \tilde{Z}(\beta).$$

The first inequality follows from the observation that $\tau = -1$ corresponds exactly to $\tilde{Z}(\beta)$. Since the configurations σ and σ_τ obtained by replacing the spins on A with τ differ in at most $4|A|$ edges, and thus $H(\sigma) - 4|A| \leq H(\sigma_\tau) \leq H(\sigma) + 4|A|$. \blacksquare

Let $\tilde{m}(\sigma)$ be the magnetization on the uncorrupted set of vertices. Define $\tilde{S} = \{\sigma : \tilde{m}(\sigma) > 0\}$. Note that this represents half of the possible configurations, and since the corrupted vertices have spins fixed to -1 , it follows that $\tilde{\mu}(\tilde{S}) < \frac{1}{2}$. Note that $\tilde{Q}(\tilde{S}, \tilde{S}^c) \leq \tilde{\mu}(\{\sigma : \tilde{m}(\sigma) = 0\})$, so we proceed by bounding the right hand side. Note that $\tilde{m}(\sigma) = 0$ implies $m(\sigma) = -|A|$ with the assumption of -1 spins on the corrupted set. Thus by the above cited theorem,

$$\begin{aligned} \tilde{Z}(\beta) \tilde{\mu}(\{\sigma : \tilde{m}(\sigma) = 0\}) &= \sum_{\tilde{m}(\sigma)=0} e^{-\beta H(\sigma)} \leq \sum_{m(\sigma)=-|A|} e^{-\beta H(\sigma)} \\ &= Z(\beta) \mu(\{\sigma : m(\sigma) = -|A|\}) \leq Z(\beta) e^{-c_\delta \beta(2n+1)} \end{aligned}$$

for some $\delta > 0$ and n sufficiently large. It follows from Lemma 35 that

$$\tilde{\mu}(\{\sigma : \tilde{m}(\sigma) = 0\}) \leq \frac{Z(\beta)}{\tilde{Z}(\beta)} e^{-c_\delta \beta(2n+1)} = e^{-c_\delta \beta(2n+1)+o(n)}.$$

To conclude, it suffices to show that $\tilde{\mu}(\tilde{S})$ is not too small. Note that $m(\sigma) > |A|$ implies $\tilde{m}(\sigma) > 0$, as the contribution from the corrupted set to the magnetization is bounded by $-|A|$. Thus,

$$\begin{aligned} \tilde{Z}(\beta) \tilde{\mu}(\tilde{S}) &= \sum_{\sigma \in \tilde{S}} e^{-\beta H(\sigma)} \geq \frac{1}{2^{|A|} e^{o(n)}} \sum_{m(\sigma) > |A|} e^{-\beta H(\sigma)} \\ &= \frac{Z(\beta)}{2^{|A|} e^{o(n)}} \mu(\{\sigma : m(\sigma) > |A|\}) \\ &\geq \frac{Z(\beta)}{2^{|A|} e^{o(n)}} \left(\frac{1}{2} - |A| e^{-c_A \beta(2n+1)} \right) \geq Z(\beta) e^{-o(n)} \end{aligned}$$

Here, the penultimate inequality follows from the cited theorem and the symmetry of the Ising model, while the final inequality holds for n sufficiently large. Again by Lemma 35, we can deduce that $\tilde{\mu}(\tilde{S}) \geq e^{-o(n)}$. Thus,

$$\tilde{\Phi}^* \leq \tilde{\Phi}(\tilde{S}) \leq \frac{e^{-c_A\beta(2n+1)+o(n)}}{e^{-o(n)}} = e^{-c_A\beta(2n+1)+o(n)}.$$

■

Appendix H. Proof of Proposition 26

Proof Fix a vertex $v \in V \setminus A$. Note that the effect of vertices at distance k from v is of order β^k . In particular, if we condition on the second neighborhood of v , then $\mu(\sigma_v = +1) = \frac{1}{2} \pm O((\beta d)^2)$. Thus, if we impose that the spins on A are fixed to $+1$, then

$$\mu(\sigma_v = +1 | \sigma_A = +1) = \frac{1}{2} + \beta |N(v) \cap A| + \beta^2 |N(N(v)) \cap A| + O(\beta^3 d^3).$$

By linearity of expectation,

$$\begin{aligned} \mathbb{E}[\{v : \sigma(v) = +1\} | \sigma_A = +1] &= m + \sum_{v \in V \setminus A} \mu(\sigma_v = +1 | \sigma_A = +1) \\ &= \frac{n+m}{2} + \beta \sum_{a \in A} d_{V \setminus A}(a) + O(m(\beta d)^2) \end{aligned}$$

where the last equality follows from the maximum degree assumption and summing the resulting geometric series. For our choice of parameters, this lives on the interval

$$\left[\frac{n+m}{2} + \beta \sum_{a \in A} d_{V \setminus A}(a) - \frac{\beta}{4}, \frac{n+m}{2} + \beta \sum_{a \in A} d_{V \setminus A}(a) + \frac{\beta}{4} \right].$$

For $t \geq M(\epsilon)$, we know that the total variation distance between μ and the distribution of X_t^A is at most ϵ . Since $\left| \{v : \tilde{X}_t^A(v) = +1\} \right|$ is a random variable bounded by n , we know that

$$\left| \mathbb{E} \left[\left| \{v : \tilde{X}_t^A(v) = +1\} \right| \right] - \mathbb{E} \left[\left| \{v : \sigma_v = +1 | \sigma_A = +1\} \right| \right] \right| \leq \epsilon n = \frac{\beta}{4}.$$

By the computations above, the triangle inequality implies that

$$\left| \mathbb{E} \left[\left| \{v : \tilde{X}_t^A(v) = +1\} \right| \right] - \left(\frac{n+m}{2} + \beta \sum_{a \in A} d_{V \setminus A}(a) \right) \right| \leq \frac{\beta}{2}.$$

It follows from this that in order to maximize $\mathbb{E} \left[\left| \{v : \tilde{X}_t^A(v) = +1\} \right| \right]$ it suffices to maximize the integer $\sum_{a \in A} d_{V \setminus A}(a)$. ■

Appendix I. Proof of Lemma 32

Proof Note that $f_v(S) = \frac{1}{2} [1 + \tanh(\beta(2|S| - d_v + h))]$. Thus, for submodularity, it suffices to prove that

$$\begin{aligned} \tanh(\beta(2|S| - d_v + h)) + \tanh(\beta(2|T| - d_v + h)) \\ \geq \tanh(\beta(2|S \cup T| - d_v + h)) + \tanh(\beta(2|S \cap T| - d_v + h)). \end{aligned}$$

We can compute that

$$\tanh(x) + \tanh(y) = \frac{e^x - e^{-x}}{e^x + e^{-x}} + \frac{e^y - e^{-y}}{e^y + e^{-y}} = 2 \cdot \frac{e^{x+y} - e^{-x-y}}{e^{x+y} + e^{x-y} + e^{-x+y} + e^{-x-y}}. \quad (1)$$

It holds that $|S| + |T| = |S \cup T| + |S \cap T|$ so we only need to compare $\exp(|S| - |T|) + \exp(-|S| + |T|)$ with $\exp(|S \cup T| - |S \cap T|) + \exp(-|S \cup T| + |S \cap T|)$. The function $e^x + e^{-x}$ is increasing for $x > 0$, and $|S| - |T| \leq |S \cup T| - |S \cap T|$ so

$$\exp(|S| - |T|) + \exp(-|S| + |T|) \leq \exp(|S \cup T| - |S \cap T|) + \exp(-|S \cup T| + |S \cap T|).$$

Substituting this into (1) and noting that $2|S| + 2|T| - 2d_v + 2h \geq 0$ implies the desired inequality. \blacksquare

Appendix J. Proof of Theorem 33

Proof Let $A_t = \{v \in V : \sigma_t^{S \cap T}(v) = +1\}$, $B_t = \{v \in V : \sigma_t^S(v) = +1\}$, $C_t = \{v \in V : \sigma_t^T(v) = +1\}$, and $D_t = \{v \in V : \sigma_t^{S \cup T}(v) = +1\}$. We construct a coupling such that $A_t \subseteq B_t \cap C_t$ and $D_t \subseteq B_t \cup C_t$.

At each time step t , draw $v_t \sim \text{Unif}(V)$ and $\theta_t \sim \text{Unif}([0, 1])$. We define the following notation for the set of neighbors of v_t with +1 spin in each chain:

$$\begin{aligned} A(v_t) &:= N(v_t) \cap (A_{t-1} \cup (S \cap T)) \\ B(v_t) &:= N(v_t) \cap (B_{t-1} \cup S) \\ C(v_t) &:= N(v_t) \cap (C_{t-1} \cup T) \\ D(v_t) &:= N(v_t) \cap (D_{t-1} \cup S \cup T) \end{aligned}$$

Apply the following update rules to the respective chains:

$$\begin{cases} \sigma_t^{S \cap T}(v_t) = +1 & \text{if } \theta_t < f_{v_t}(A(v_t)) \text{ or } v \in S \cap T \\ \sigma_t^S(v_t) = +1 & \text{if } \theta_t < f_{v_t}(B(v_t)) \text{ or } v \in S \\ \sigma_t^T(v_t) = +1 & \text{if } \theta_t < f_{v_t}(A(v_t)) \text{ or } 1 - \theta_t < f_{v_t}(C(v_t)) - f_{v_t}(A(v_t)) \text{ or } v \in T \\ \sigma_t^{S \cup T}(v_t) = +1 & \text{if } \theta_t < f_{v_t}(B(v_t)) \text{ or } 1 - \theta_t < f_{v_t}(D(v_t)) - f_{v_t}(B(v_t)) \text{ or } v \in S \cup T \end{cases}$$

and the spins are updated to -1 otherwise. Since $A_{t-1} \subseteq B_{t-1}$ by monotonicity $f_{v_t}(A(v_t)) \leq f_{v_t}(B(v_t))$, we have that $v_t \in A_t$ only if $v_t \in B_t$ and $v_t \in C_t$. Moreover if v_t is removed from either B_t or C_t in this update, then it will be removed from A_t as well. Thus, $A_t \subseteq B_t \cap C_t$. If

$v_t \in D_t$, then either $\theta_t < f_{v_t}(B(v_t))$ in which case $v_t \in B_t$ or $1 - \theta_t < f_{v_t}(D(v_t)) - f_{v_t}(B(v_t))$, in which case $v_t \in C_t$ since

$$f_{v_t}(D(v_t)) - f_{v_t}(B(v_t)) < f_{v_t}(C(v_t)) - f_{v_t}(A(v_t)).$$

Indeed

$$\begin{aligned} f_{v_t}(B(v_t)) + f_{v_t}(C(v_t)) &\geq f_{v_t}(B(v_t) \cap C(v_t)) + f_{v_t}(B(v_t) \cup C(v_t)) \\ &\geq f_{v_t}(A(v_t)) + f_{v_t}(D(v_t)) \end{aligned}$$

where we used submodularity in the first inequality, and monotonicity and the induction hypothesis in the second inequality. Similarly if v_t is removed from both B_t and C_t then it will be removed from D_t as well, so $D_t \subseteq B_t \cup C_t$.

Now,

$$|B_t| + |C_t| = |B_t \cap C_t| + |B_t \cup C_t| \geq |A_t| + |D_t|$$

Taking expectations shows that $\mathbb{E} \left[\left| \{v : \tilde{X}_t^S(v) = +1\} \right| \right]$ is submodular. ■