

Hypergraph coloring games and voter models

Fan Chung Alexander Tsiatas
Department of Computer Science and Engineering
University of California, San Diego
{fan,atsiatas}@cs.ucsd.edu

Abstract

We analyze a network coloring game on hypergraphs which can also describe a voter model. Each node represents a voter and is colored according to its preferred candidate (or undecided). Each hyperedge is a subset of voters that can interact and influence one another. In each round of the game, one hyperedge is chosen randomly, and the voters in the hyperedge can change their colors according to some prescribed probability distribution. We analyze this *interaction model* based on random walks on the associated weighted, directed state graph. We consider three scenarios — a memoryless game, a partially memoryless game and the general game using the memoryless game for comparison and analysis. Under certain ‘memoryless’ restrictions, we can use semigroup spectral methods to explicitly determine the spectrum of the state graph, and the random walk on the state graph converges to its stationary distribution in $O(m \log n)$ steps, where n is the number of voters and m is the number of hyperedges. This can then be used to determine an appropriate cut-off time for voting: we can estimate probabilities that events occur within an error bound of ϵ by simulating the voting game for $O(\log(1/\epsilon)m \log n)$ rounds. Next, we consider a partially memoryless game whose associated random walk can be written as a linear combination of a memoryless random walk and another given random walk. In such a setting, we provide bounds on the convergence time to the stationary distribution, highlighting a tradeoff between the proportion of memorylessness and the time required. To analyze the general interaction model, we will first construct a companion memoryless process and then choose an appropriate *damping constant* β to build a partially memoryless process. The partially memoryless process can serve as an approximation of the actual interaction dynamics for determining the cut-off time if the damping constant is appropriately chosen either by using simulation or depending on the rules of interaction.

1 Introduction

We consider a network coloring game, motivated by human behavioral experiments [14, 16] conducted in a network setting. The network coloring game can be formulated as the following interaction-based voter model:

A set of voters is modeled as the vertex set of a hypergraph $H = (V, E)$. Each hyperedge $g \in E$ represents a small, possibly overlapping group representing social interactions and discussions (such as lunchtime hallway or office discussions, blog commentary, television viewership, and Web forums). At time $t = 0$, each voter has a *color* representing an initial preference or is undecided. Then at each time step, one hyperedge g is randomly selected according to some prescribed probability distribution on E . After the voters in g interact with one another, the voters’ preferences can change probabilistically. This process is repeated for many rounds, and the coloring configuration of the voters evolves.

Many natural questions arise. Will the coloring configurations converge under certain conditions? If the coloring patterns diverge or oscillate, what would be an appropriate time frame to stop the model? When the model is run for some prescribed number of rounds, how does the observed coloring configuration behave? Can the observed coloring configuration be something other than ‘random’? What is the stationary distribution of the observed coloring configuration, if it exists?

The behavioral experiments of Kearns et al. [11, 12, 14, 15] were the original motivation for our model. In these simulations, actual human agents were given a common graph coloring task with varying monetary rewards for their timely completion. This is a strategic component that is not necessarily built into our interaction-based model, but the probability distributions for color change can potentially be derived using game-theoretic analysis. The experiments allowed agents differing “views” of information about the specific graph problem structure. Thus, agents only had limited information and could only make decisions based on the color configurations of smaller subsets of nodes (but not just pairs). This gives a compelling reason to model the voters as nodes of a hypergraph and not a traditional graph where edges are limited to two endpoints.

These problems have been quite elusive in spite of extensive study either by simulation [11, 12, 14, 15] or by using various other proposed voter models [7, 10, 18, 19] for various types of networks [23]. The results vary widely with the specific set of rules or dynamics and point to evidence of intrinsic difficulties involved in this general voting model.

Although the general voting and coloring problems may be intractable to analyze tightly, we will consider several special cases that will help in evaluating the more general case. The first is a special case of *memoryless* voter models. In a memoryless voter model, the voters do not consider their current preferences when formulating new ones. Under these assumptions, we can show that the dynamics of coloring configurations can be analyzed quite precisely in several aspects.

Our coloring game is played on a hypergraph where each hyperedge is formed by a group to model multi-party simultaneous interaction. Although the coloring configurations of the nodes in the hypergraph do not converge in general, we can model the rounds of coloring games by conducting random walks on the associated directed state graph which contains coloring configurations as nodes. It turns out, for memoryless interactions, the spectrum of the state graph has elegant expressions. Using the formula for the eigenvalues, we can determine the rate of convergence for random walks on the state graph, which can then be used for determining a cut-off time for the voting game.

The eigenvectors of the random walk on the state graph contain rich information. In particular, the stationary distribution, which is the eigenvector associated with eigenvalue 1, tells us the distribution of the observed coloring configuration after the cut-off time. In contrast with undirected graphs, the stationary distribution a random walk on a directed graph does not always exist and is not easy to determine, especially for state graphs with exponentially large size. Even for simple hypergraphs H such as the path or cycle, the stationary distribution for random walks on the state graph have complicated forms quite different from the uniform distribution [6]. Nevertheless, we will show that for events such as “red wins by at least 5% within an error bound ϵ ” can be determined by simulating the voter model for $O(\log(1/\epsilon)m \log n)$ rounds, where m is the number of hyperedges in H , provided that the hyperedges are always chosen uniformly at random, and the interactions are memoryless.

Next, we proceed to consider the partially memoryless model for which, with some probability β , the interaction process is memoryless and with probability $1 - \beta$, the process is not required to be memoryless. Specifically, a partially memoryless voter model is one whose associated state-graph

random walk can be decomposed into two parts: $P = \beta M + (1 - \beta)P'$, where M is a memoryless random walk and P' is any given random walk. Using results from the memoryless case and also from general directed Laplacians [5], we can analyze partially memoryless voter games as well.

The memoryless condition might seem to be too strict to be satisfied by real-world social interactions, but it can still be of interest to serve as a type of benchmark for the sake of comparison. In particular, the partially memoryless case provides a natural framework for evaluating a more general interaction voter process. For a given interaction voter process with its associated random walk P , we can construct a corresponding memoryless process $P' = \beta M + (1 - \beta)P$ where M is a memoryless process. The determination of M depends on β and simulation of P with the cut-off time depending on β . The details for choosing β and M will be given later in Section 7.

Related work

The coordination or consensus game has been extensively studied in evolutionary network game theory [8, 13, 25]. In a set of controlled behavioral experiments, a simple voting game [12] as well as a biased version [14] were simulated on a set of small, constructed topologies. The rules varied, but involved small financial rewards for success and differing amounts of information visibility. The results showed that consensus was often reached within a certain timeframe, though there were many cases of failure as well. It is desirable to further explain these social network phenomena.

One method to analyze various consensus games is to use the standard voter model [7, 10], where one agent is selected randomly during each round, assuming the vote of a randomly-selected neighbor. Some empirical work has shown that its performance depends on network topology [22]. Several proofs have been given [1, 24] using coalescing random walks to show that the expected time to reach consensus under the voter model is $O(n^3 \log n)$, where n is the number of voters in the network. For the biased voting game, this method takes exponential time [16], but it can be used as a subroutine for an algorithm converging in $O(n^8 \log n)$ expected time. However, elections in practice rarely end with a unanimous decision, and our model reflects this truth.

Other recent work consider alternative dynamics for the voting game, including Glauber dynamics [18] and pieces of “advice” [19] given to the voters. These models also result in consensus, and the time required is related to graph-theoretic quantities such as the tilted cutwidth, diameter, and broadcast time, often difficult to compute or reason about. Some of these models assume that the voters know approximate values for these quantities, which our model avoids. These models again result in unanimous coordination, whereas our proposed model is more likely to result in more realistic voting configurations.

The experiments of Kearns et al. [12, 14] show that consensus is indeed not always reached, especially for biased voting games. But not much has been said thus far about the distribution of voting configurations should a true consensus be unattainable.

Furthermore, nearly all work done on the network coordination game has focused on pairwise relationships between voters. But voters often interact in groups larger than two, and most of the models described so far have not taken this possibility into account. Such interactions as town-hall debate, reader comments on news articles, and dinner-table discussions can all possibly be considered in the framework of our hypergraph model.

A summary of our results

We consider a variation on the voter model that takes into account multi-voter interactions for the coordination game. Instead of limiting ourselves to pairwise relationships between voters, we model the interaction network as a hypergraph $H = (V, E)$, where each of n vertices is a voter and each of m hyperedges represents a group that can interact.

Our interaction model can be described as follows. At the beginning, all voters have some initial views, voting for one of several candidates, or starting undecided. In each round, one hyperedge g is selected randomly, and an *interaction* $X_{g,\tau}$ takes place: the randomly selected hyperedge g changes its pattern to τ . In the most general case, the probability of $X_{g,\tau}$ occurring can depend on the current coloring configuration σ of H . We denote by $p(g, \sigma, \tau)$ the probability of $X_{g,\tau}$ occurring when σ is the current coloring configuration on H , and for all σ , it must be the case that $\sum_{g,\tau} p(g, \sigma, \tau) = 1$.

If the interactions are memoryless, then the probability of $X_{g,\tau}$ occurring is constant across all voting configurations σ on H . In this case, we can omit σ and denote by $p(g, \tau)$ the probability that $X_{g,\tau}$ occurs. When this probability only has two parameters, it can be assumed that the interaction is memoryless. The model is then simulated for some preset number of rounds.

While the game is taking place, the state can be described as a *voting configuration* or *coloring configuration* of the voters among the candidates or undecided. If there are r possible votes (including undecided), there are r^n possible configurations, and we can construct a *state graph* H^* , where the vertices are coloring configurations, and a directed edge connects u to v if the state v is reachable from u in one round of the interaction model. Thus, a simulation of the coloring game using the interaction model can be completely described as a random walk on the state graph H^* .

Let an *event* $A \subset V(H^*)$ be a subset of the states in the state graph. These events can represent numerous scenarios: for example, the states where more than half of voters choose red, or the states which have a specific trend over some subsets of the voters, representing a voting district or municipality. In general, A can be any event of interest. We will show that for memoryless interactions, we can estimate $\Pr[A]$ within a probabilistic error bound by using a sufficient number of samplings, as long as A contains enough states.

Theorem 1. *Let A be an event or subset of all possible coloring configurations on a hypergraph H on n vertices and m edges. Suppose the interactions are memoryless, the probability of selecting a hyperedge g in any given round is at least $1/(\alpha m)$ for some α , and for any set $S \subseteq V(H)$ there are at least $|S|$ edges incident to vertices in H . The probability $\Pr[A]$ that A occurs at any cut-off time after $O(\alpha m \log n)$ rounds of simulation can be estimated within an error bound of ϵ using $O(\alpha \log(1/\epsilon) m \log(n) / \Pr[A])$ rounds of simulation.*

Note that for the special case that the hyperedges are chosen uniformly at random, the number of rounds of simulation for convergence is:

$$O(\log(1/\epsilon) m \log(n) / \Pr[A]).$$

The main tools that we use to prove the above theorem are sampling and fast mixing of random walks on the state graph H^* associated with memoryless strategies. In particular, the spectrum of such random walks on the state graph H^* can be determined by using spectral techniques originating in the analysis of card shuffling [3, 4], self-organizing search [9], hyperplane arrangement [2] and semigroup random walks as well as the recent work on edge flipping games in graphs [6].

For the partially memoryless case, where the random walk P on the state graph H^* can be written as $P = \beta M + (1 - \beta)P'$ with memoryless M , we have the following result showing a trade-off between the convergence time and β :

Theorem 2. *Let A be an event or subset of all possible coloring configurations on a hypergraph H on n vertices and m edges. Suppose the interactions are partially memoryless with parameter β , the probability of selecting a hyperedge g in any given round is at least $1/(\alpha m)$ for some α , and for any set $S \subseteq V(H)$ there are at least $|S|$ edges incident to vertices in H . The probability $\Pr[A]$ that A occurs at any cut-off time after $O(\alpha m \log n)$ rounds of simulation can be estimated within an error bound of ϵ using $O\left(\frac{\alpha \log \frac{1}{\epsilon} m \log n}{\beta \Pr[A]}\right)$ rounds of simulation.*

Finally, for the most general case, where the interactions are not memoryless at all, it is well known to be very difficult to deal with [6]. This is because the associated random walk can have exponentially small eigenvalues, and convergence time can vary widely based on the specific dynamics. Nevertheless, we will show how to use insight from the memoryless and partially memoryless models to reason about the general voter model. By choosing an appropriate damping constant β and extracting the memoryless version of a given process, we can then construct a partially memoryless process which can be used to approximate the given interactive process.

2 The voting game on a hypergraph as a random walk on the associated state graph

We first describe the interaction model as a coloring game on a hypergraph $H = (V, E)$. The set V of nodes consists of all voters and each hyperedge $g \in E$ represents a group of voters who can interact with one another. For a hyperedge g and coloring pattern τ on the voters in g , we let the *interaction* $X_{g,\tau}$ denote one step of the model that moves a coloring configuration of V to another by changing the color pattern to τ for voters in g .

The state graph H^* is a weighted directed graph whose vertices are coloring configurations of $V(H)$. To distinguish from nodes in H , we sometimes call a vertex in $V(H^*)$ a *state*. There is a directed edge from a state u in H^* to another state v if there is a hyperedge g and a coloring pattern τ on g such that the interaction $X_{g,\tau}$ moves u to v . The weight on the edge (u, v) is determined by the probability that $X_{g,\tau}$ occurs during state u . We denote by $p(g, u, \tau)$ the probability that, in state u , hyperedge g is selected and changes its color configuration to τ .

Note that for any directed graph with weighted edges, we can define a typical random walk with a transition from u to v occurring with probability $w(u, v) / \sum_z w(u, z)$. Therefore, the typical random walk associated with the state graph H^* simulates the evolving configurations in the voter interaction game.

Starting from a coloring configuration u , a sequence of interactions

$$X_{g_1, \tau_1} X_{g_2, \tau_2} \dots X_{g_t, \tau_t}$$

induces a series of changes in the coloring configuration and can then be viewed as a walk starting from u , traversing t directed edges on the state graph H^* . This correspondence leads to the following lemma whose proof follows from the Perron-Frobenius Theorem.

Lemma 1. *The voter interaction game with interactions as described above does not converge to an equilibrium in general. Instead, from any initial coloring configuration, the resulting configuration*

after t rounds of simulation, is s with probability approaching $\pi(s)$, where π is the stationary distribution of the random walk on the state graph H^* , as long as t is sufficiently large and H^* is strongly connected and aperiodic.

An interaction $X_{g,\tau}$ is said to be *nontrivial* if the associated probability $p(g,\sigma,\tau)$ is nonzero. In the special case that all nontrivial strategies are consistent with some coloring pattern τ (for example, all red), then the coloring configuration will reach an equilibrium. (Note that this is a special case where the interactions are memoryless.) Suppose there is a state s for which all the nontrivial interactions are of the form X_{g,s_g} , where s_g is simply the coloring configuration of s restricted to nodes in g . If this is the case, then starting from any initial configuration, the voting game will converge using standard coupon-collector probabilistic arguments:

Lemma 2. *In the voter interaction game, suppose there exists a coloring configuration s such that all the nontrivial interactions are of the form X_{g,s_g} , where s_g denotes the coloring pattern of s restricted to voters in g . Starting from any initial configuration, the voting game converges to s after t rounds of simulation with probability at least $1 - e^{-c}$ if*

$$t \geq \frac{\log n + c}{\min_{v \in V} \sum_{\substack{g \in E \\ g \ni v}} p(g, s_g)},$$

where n is the number of voters and $p(g, s_g)$ is the probability associated with the interaction. In the case that every vertex is incident to exactly d hyperedges and each hyperedge is chosen with equal probability, the above inequality is just $t \geq n(\log n + c)$.

For the remainder of this paper, we will assume that the given hypergraph and interaction dynamics yield a state graph H^* that is aperiodic and strongly connected. We will refer to the stationary distribution π accordingly.

We remark that Lemma 1 reduces the voter interaction game to random walks, and the rate of convergence depends on the eigenvalues of the directed state graph. These values can be complex, and in the most general case, the spectral gap can be exponentially small. Nevertheless, we will consider memoryless interactions which allow us to have real eigenvalues for the state graph, and we can use these techniques to derive some bounds for partially memoryless interactions as well.

3 Memoryless interactions and semigroup spectral graph theory

For a random walk on the state graph H^* , we can describe a path in terms of the interactions $\{X_{g,\tau}\}$ that take place to follow the path. Thus, it is convenient to describe random walks as sequences of interactions. For a sequence $S = X_{g_1,\tau_1} X_{g_2,\tau_2} \dots X_{g_t,\tau_t}$ and a state u , we say $S = u$ if the interaction game ends up in state u after following the path described by S . For two sequences, we say $S_1 = S_2$ if both paths end at the same state.

We say the nontrivial interactions $\{X_{g,\tau}\}$ are *memoryless* if the probability of performing it does not depend on the current state. An equivalent definition is that for memoryless interactions, a repeated interaction $X_{g,\tau}$ means that earlier occurrences have no effect. Namely, for any three sequences of interactions S_1, S_2, S_3 of any length, then

$$S_1 X_{g,\tau} S_2 X_{g,\tau} S_3 = S_1 X_{g,\tau} S_2 S_3.$$

If the interaction strategies are memoryless, we can view them as members of a special type of semigroup known as a *left-regular band* or LRB, first studied in the 1940's [17, 21]. An LRB is a semigroup where every element is idempotent, and for any two elements $x, y \in S$, $xyx = xy$. We define the product of two interactions $X_{g,\tau}$ and $X_{g',\tau'}$ to be the two interactions in sequence. If the interactions are memoryless, it is easy to see that the semigroup S generated by all nontrivial interactions $X_{g,\tau}$ is a LRB. This allows us to apply techniques in [2, 3, 4, 6] to the voter interaction game. In particular, the associated random walk on the state graph for memoryless interactions has a clean form:

Theorem 3. *Suppose that the voter interaction game on a hypergraph $H = (V, E)$ in r colors has memoryless interactions $X_{g,\tau}$ for $g \in E$ and coloring patterns τ on g . If $p(g, \tau)$ is the probability of choosing g and coloring voters in g with the coloring pattern τ , then the random walk on the associated state graph H^* has an eigenvalue λ_T for every subset $T \subseteq V$:*

$$\lambda_T = \sum_{\substack{g,\tau \\ g \subseteq T}} p(g, \tau)$$

with multiplicity $(r - 1)^{n-|T|}$.

For the specific case where the hyperedges are selected uniformly at random and there are only two colors, we note that the eigenvalues have an even cleaner form: for each subset $T \subseteq V$, there is an eigenvalue:

$$\lambda_T = \frac{|\{g \in E | g \subseteq T\}|}{m}$$

with multiplicity 1.

To prove Theorem 3, we need to explore further properties of LRB semigroups. The proof of the corollary follows from Theorem A in [6], and the theorem follows by generalizing the techniques to r colors. But in order to use these results, we must first interpret semigroup terminology in terms of the voter interaction game. These details appear in the appendix, and further details about the terminology can be found in [3].

4 The cut-off time for voter interaction games

Our methods also address some of the questions that arise in the recent human network experiments of Kearns et al. [16]. The voters are given a hard deadline, often arbitrarily set by various entities without justification. Furthermore, for different stopping times, the outcome of the network experiments varied widely from consensus to chaos. Using our interaction model and spectral techniques, we will prove the following theorems about the interaction model's convergence properties, as well as a mathematical interpretation of the resulting voting configuration after convergence is reached.

For voter interaction games with memoryless interactions, we have the following theorem. The proof follows by using the spectrum of H^* (derived in the previous section) to bound the total variation distance between the random walk and its stationary distribution after t steps:

Theorem 4. *Suppose the interaction model is simulated on a hypergraph $H = (V, E)$ with $|V| = n$, and each voter in V is colored with one of r colors. If the interactions are memoryless, then the*

total variation distance between the random walk on the state graph H^* denoted by the transition probability matrix P after t rounds of simulation and its stationary distribution π is given by

$$\begin{aligned} \|P^t - \pi\|_{TV} &= \max_{A \subseteq V} \max_y \left| \sum_{x \in A} P^t(y, x) - \pi(x) \right| \\ &\leq \sum_{T \subseteq H} \lambda_T^t (r-1)^{n-|T|}. \end{aligned}$$

We remark that the above bound can be somewhat improved by restricting T to be the co-maximal subsets, although asymptotically the bound is still the same. Using this bound on the total variation distance, we can derive the convergence time for the interaction model:

Theorem 5. *On a hypergraph $H = (V, E)$ with $|V| = n$, $|E| = m$, suppose the voter interaction model is simulated with memoryless interactions, and the probability of choosing a hyperedge g in any round is uniform over E . Suppose for any set S of k voters, there are at least k hyperedges involving voters in S . The random walk on the state graph H^* converges to its stationary distribution in $O(m \log n)$ steps.*

To prove Theorem 5, we use the derived spectrum from Theorem 3 in our derived bound on the total variation distance from Theorem 4. The algebraic details are left for the appendix.

Theorems 4 and 5 imply a method for choosing a stopping point for the interaction model: enough time for a desired level of convergence to a stationary distribution π . This also indicates what the voting configuration among the agents looks like at any time after it has converged: the votes are a random sample from π from all the voting configurations in the state graph.

5 Estimating the expected value of a given event

Although estimating individual components of π can be computationally intractable, for memoryless interactions, we can effectively use sampling to estimate the probability of an event A , as long as the event has enough probability mass. For a general π , it can be difficult to reason about its components, since it contains exponentially many elements, corresponding to the state graph. Even for simple graphs such as the path of length k , the exact stationary distribution on the state graph is quite complex [6]. Additionally, because most of the components of π are exponentially small, even estimating π can be quite difficult. But in practice, the exact stationary distribution is not of utmost importance. Instead, it is much more revealing and tractable to reason about larger events that capture a larger portion of π . We will use the following fact:

Theorem 6. ([20]). *Let A be an event and $\pi(A) = \sum_{s \in A} \pi(s)$ be the probability that the outcome is in A . Let $\delta, \epsilon \in (0, 1)$. Suppose that after N samplings, X is the proportion of times the outcome was in A . Then*

$$\Pr[(1 - \delta)\pi(A) \leq X \leq (1 + \delta)\pi(A)] \leq 1 - \epsilon,$$

as long as $N \geq O\left(\frac{\log(1/\epsilon)}{\pi(A)c(\delta)}\right)$, where $c(\delta)$ only depends on δ .

We can use this to prove Theorem 1. Here, we consider the case where each hyperedge is chosen with uniform probability $1/m$, but the same argument will hold using the looser $1/(\alpha m)$ bound.

Proof. (for Theorem 1) For an event A , suppose we are given an initial state f which we denote as a row vector indexed by states in $V(H^*)$. For any integer t , the coloring configuration we observe after t rounds of the voter interaction model is in A with probability

$$\mathbb{E}_t[fA] = \sum_{x \in A} f P^t(x)$$

where P denotes the transition probability matrix of the random walk on H^* .

By combining Theorems 3 and 4, we have

$$\begin{aligned} |\mathbb{E}_t[A] - \pi(A)| &\leq \max_y \left| \sum_{x \in A} P^t(y, x) - \pi(x) \right| \\ &= \|P^t - \pi\|_{TV} \\ &\leq \sum_{T \subseteq V} \lambda_T^t (r-1)^{n-|T|}. \end{aligned}$$

For the case where each hyperedge is chosen with probability $1/m$, where m is the total number of hyperedges, we have

$$\begin{aligned} |\mathbb{E}_t[A] - \pi(A)| &\leq n^2 \left(1 - \frac{1}{m}\right)^t \\ &\leq \epsilon \end{aligned}$$

if $t > 2m(\log n + \log(1/\epsilon))$.

Now we use Theorem 6, by breaking up the voter interaction game into N phases where $N = O\left(\frac{\log(1/\epsilon)}{\pi(A)c(\delta)}\right)$ and each phase consists of $t = O(m \log n)$ rounds. The proportion of phases where the outcome is in A satisfies:

$$\Pr[(1 - \delta)\pi(A) < |X - \pi(A)| < (1 + \delta)\pi(A)] \geq 1 - 2\epsilon. \quad (1)$$

□

6 The interaction model with partially memoryless interactions

We say that a set of interactions is *partially memoryless* if the random walk on the state graph H^* can be decomposed into two parts, one of which is memoryless. Specifically, if the random walk transition matrix is P , there is a $\beta \in (0, 1)$ such that we can write

$$P = \beta M + (1 - \beta)P',$$

where the interactions described by M are memoryless and P' is another transition probability matrix without any restriction.

One way a partially memoryless interaction model can arise is if at each step, the voters in the selected hyperedge interact memorylessly with some probability β , allowing their actions to depend on the current step with probability $1 - \beta$. It is also possible that the transition probability matrix P' is not explicitly built this way, but it can nevertheless be expressed as a partially memoryless model. In this sense, the parameter β can be viewed as describing how memoryless the model is, and its properties will depend on β :

Theorem 7. *On a hypergraph $H = (V, E)$ with $|V| = n, |E| = m$, suppose the voter interaction model is simulated with partially memoryless interactions, and the probability of choosing a hyperedge g in any round is uniform over E . Suppose for any set S of k voters, there are at least k hyperedges involving voters in S . The random walk on the state graph H^* converges to its stationary distribution in $O(\frac{m \log n}{\beta})$ steps.*

The partially memoryless structure allows us to use the results from the fully memoryless case, with an additional factor of $1/\beta$ when analyzing the spectrum. The details of the proof for Theorem 7 are left for the appendix. With this result, we can prove Theorem 2. The proof follows by using the same techniques as the proof for Theorem 1; we use Theorem 6 in the same manner. We have the same bound as (1), but the number of rounds required in each phase is now given as in Theorem 7 (Eq. (2)).

7 The general interaction model

Thus far, we have written about the interaction model with memoryless and partially memoryless interactions. The general version of the interaction model concerns interactions that all can depend on the current and previous states, where the random walk on the state graph H^* can be quite difficult to analyze. The eigenvalues of the random walk can be real or complex, and exponentially small [6]. With $O(2^n)$ states, computing the spectrum explicitly is too expensive or infeasible. Nevertheless, we can use the previous memoryless and partially memoryless models as a basis for comparison between a specific set of state-dependent interactions and memorylessness in general.

In particular, for any general interaction voter game, we can build a companion game which is partially memoryless with one scalar parameter β . Let M denote the transition probability matrix of the random walk on H^* . We wish to construct a partially memoryless model

$$P' = \beta M + (1 - \beta)P$$

where M is memoryless and can be constructed from P and β as follows: Instead of simulating blindly, we will use Theorem 2 which provides an upper bound on the convergence time to the stationary distribution for partially memoryless interactions. If the companion model serves as an approximation for the actual voter game, it is enough to simulate the more general model $O(\frac{m}{\beta} \log n)$ time. To construct M , we collect a sample probability distribution: when a hyperedge g is selected, keep track of the resulting coloring configurations. Then we can use these samples to build a partially memoryless model: first select a hyperedge g , then with probability β , randomly select one of the collected sample coloring configurations on g . Note that once the samples have been collected, this step is memoryless. With probability $1 - \beta$, we just use the original memory-dependent dynamics.

The problem of determining the appropriate value of β can be quite difficult since the required length of time until convergence may be exponential. A feasible heuristic approach is a combination of a series of iteration and simulation. Such process also can be used to get a sense of how “memoryless” the more general model really is. In other words, an alternative interpretation of the damping constant β is just the ratio of the rate of convergence of the memoryless random walk and the actual random walk. In general, the value of β can range from 1 to some exponentially small values, reflecting the fact that some models have higher extent of memorylessness than others. So even though the general model is notoriously difficult to reason about, we can still quantify its convergence time empirically by comparing with the partially memoryless model.

To illustrate this comparative process, we consider a traditional voting game on networks. We are given a graph G , and each node has a starting color. Then at each time step, a node is randomly selected, and it takes the color of one of its neighbors, selected uniformly at random. In our setting, there is a hyperedge for every node v , consisting of v and its neighbors, and whenever it is chosen, v changes color to one of its neighbors' colors. It should be clear that these dynamics are completely memory-dependent: the coloring configuration at time $t + 1$ always depends on the state at time t . But we can build a partially memoryless process using the method described in the previous paragraph.

We demonstrate this method using Zachary's karate network [26] as an example graph. This graph has 34 nodes, and we initially assign one of 9 colors randomly to each node. Two sample runs of the consensus game are shown in Figs. 1a and 1b. Each row represents one node, and the colors change as time moves from left to right.


We note that the state graph can be as large as $9^{\binom{34}{2}}$ nodes. The convergence bound for a memoryless game of similar size is about $O(m \log n)$ steps with $n = 34$ and $m = \binom{34}{2}$. This is comparable to the time limit of 1000 steps, if each step is taken to be in the range of a fraction of a second. But in our simulations, consensus may or may not be reached.

In Fig. 1c, we give an illustration of a partially memoryless version of the consensus game. Using only 100 rounds of simulation, we built the partially memoryless version and simulated it for 1000 rounds. In practice, β can be chosen empirically, by using binary search on $(0, 1)$. For Fig. 1c, we chose $\beta = 0.01$. One way to choose β is by iteratively adjusting β so that the proportion of cases that achieve consensus reaches the range of what is to be expected.

In this paper, we have used a one-parameter family of partially memoryless models to approximate a general voter game. Both the problems of finding a deterministic way to choose the parameter β and to rigorously analyze the sharpness of approximation remain open for future research.

References


- [1] D. Aldous and J.A. Fill, Reversible Markov chains and random walks on graphs. Book in preparation (1994).
- [2] T.P. Bidigare, P. Hanlon and D.N. Rockmore, A combinatorial description of the spectrum for the Tsetlin library and its generalization to hyperplane arrangements. *Duke Mathematical Journal* 99 (1999), 135–174.
- [3] K.S. Brown, Semigroups, rings, and Markov chains. *Journal of Theoretical Probability* 13 (2000), 837–938.
- [4] K.S. Brown and P. Diaconis, Random walks and hyperplane arrangements. *Annals of Probability* 26 (1998), 1813–1854.
- [5] F. Chung, Laplacians and the Cheeger inequality for directed graphs. *Annals of Combinatorics* 9 (2005), 1–19.
- [6] F. Chung and R. Graham, Edge flipping in graphs. *Advances in Applied Mathematics* 48 (2012), 37–63.



(a) One simulation of the consensus game on Zachary's karate network for 1000 steps.



(b) Another simulation of the consensus game on Zachary's karate network for 1000 steps.



(c) Simulation of the partially memoryless version of the consensus game on Zachary's karate network for 1000 steps, with $\beta = 0.01$ and 100 steps of training.

Figure 1: Simulations of the consensus game and its partially memoryless approximation.

- [7] P. Clifford and A. Sudbury, A model for spatial conflict. *Biometrika* 60 (1973), 581–588.
- [8] G. Ellison, Learning, local interaction, and coordination. *Econometrica* 61 (1993), 1047–1071.
- [9] J.A. Fill, An exact formula for the move-to-front rule for self-organizing lists. *Journal of Theoretical Probability* 9 (1996), 113–160.
- [10] R.A. Holley and T.M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model, *The Annals of Probability* 3 (1975), 643–663.
- [11] S. Judd and M. Kearns, Behavioral experiments in networked trade. *Proceedings of the 9th ACM Conference on Electronic Commerce* (2008), 150–159.
- [12] S. Judd, M. Kearns and Y. Vorobeychik, Behavioral dynamics and influence in networked coloring and consensus. *Proceedings of the National Academy of Sciences* 107 (2010), 14978–14982.
- [13] M. Kandori, G. Mailath and R. Rob, Learning, mutation, and long run equilibria in games. *Econometrica* 61 (1993), 29–56.
- [14] M. Kearns, S. Judd, T. Tan and J. Wortman, behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences* 106 (2009), 1347–1352.
- [15] M. Kearns, S. Suri and N. Montfort, An experimental study of the coloring problem on human subject networks. *Science* 313 (2006), 824–827.
- [16] M. Kearns and J. Tan, Biased voting and the Democratic primary problem. *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE 2008)*, 639–652.

- [17] F. Klein-Barmen, Über eine weitere Verallgemeinerung des Verbandsbegriffes. *Mathematische Zeitschrift* 46 (1940), 472–480.
- [18] A. Montanari and A. Saberi, Convergence to equilibrium in local interaction games and Ising models. *ACM SIGecom Exchanges* 8 (2009).
- [19] E. Mossel and G. Schoenebeck, Reaching consensus on social networks. *Proceedings of the First Symposium on Innovations in Computer Science* (ICS 2009).
- [20] R. Motwani and P. Raghavan, *Randomized Algorithms*, Cambridge University Press, 1995.
- [21] M.-P. Schützenberger, Sur certains treillis gauches. *Comptes Rendus de l’Académie des Sciences* 224 (1947), 777–778.
- [22] K. Suchecki, V. Eguíluz and M. San Miguel. Voter model dynamics in complex networks: Role of dimensionality, disorder, and degree distribution. *Physical Review E* 72 (2005), 1–8.
- [23] A. Tahbaz-Salehi and A. Jadbabaie, Consensus over ergodic stationary graph processes. *IEEE Transactions on Automatic Control* 55 (2010), 225–230.
- [24] M. Yildiz, R. Pagliari, A. Ozdaglar and A. Scaglione, Voting models in random networks. *Proceedings of the 2010 Information Theory and Applications Workshop* (ITA), 1–7.
- [25] H.P. Young, The evolution of conventions. *Econometrica* 61 (1993), 57–84.
- [26] W.W. Zachary, An information flow model for conflict and fission in small groups. *Journal of Anthropological Research* 33 (1977), 452–473.

A Semigroup details for proving Theorem 3

For our LRB semigroup S , there is a natural partial order defined by: $x \leq y \Leftrightarrow xy = y$. In other words, $x \leq y$ if the interaction or sequence of interactions represented by x is irrelevant after performing the sequence y . A semilattice $L(S)$ can be defined on S by considering the relation \preceq on S as follows: $y \preceq x \Leftrightarrow xy = x$. The equivalence class under \preceq which contains x is said to be the *support* of x , denoted by $\text{supp}(x)$, and for $x, y \in S$, $\text{supp}(xy) = \text{supp}(x) \cup \text{supp}(y)$. The support of x can be interpreted as the set of vertices whose colors were affected by the sequence of interactions given by x . The various elements of $L(S)$ are called the *flats*, and an element of S is said to be a *chamber* if its support is maximal. Therefore, the chambers of S are simply sequences of interactions x that affect the entire set of nodes V .

As given in [3, 4, 6], the eigenvalues of a random walk on chambers have an elegant form. For each flat $X \in L(S)$, there is an eigenvalue $\lambda_X = \sum_{x \in X} w_x$. Here, w_x is the probability of selecting the semigroup member x . In the voter interaction model, the chambers represent coloring configurations, and the probabilities w_x are simply the probabilities of choosing specific interactions. The flats are simply subsets of V , and it becomes clear how the eigenvalues of the random walk on H^* are derived. For the multiplicities, we note from [6] that the multiplicity of m_X of λ_X satisfies

$$\sum_{Y \succeq X} m_Y = c_X,$$

where c_Y is the cardinality of $S_{\geq Y} = S_{\geq y} = \{z \in S : z \geq y\}$, where y is any element with support Y . (The cardinality is independent of the choice of y .) It can then be seen how the multiplicities in Theorem 3 were derived.

We note that these techniques as used in [6] were developed for an edge-flipping game with two colors. The semigroup techniques used do generalize to r colors, and the full proof for Theorem 3 can be derived in that manner.

B Proof of Theorem 5

Proof. Using Theorems 3 and 4:

$$\begin{aligned} \|P^t - \pi\|_{TV} &\leq \sum_{T \subseteq H} \lambda_T^t (r-1)^{n-|T|} \\ &= \sum_{T \subseteq V} \left(\frac{|\{g \in E | g \subseteq T\}|}{m} \right)^t (r-1)^{n-|T|} \\ &\leq \sum_{k=1}^n \left(\frac{\max_{|T|=k} |\{g \in E | g \subseteq T\}|}{m} \right)^t \binom{n}{k} (r-1)^{n-k}. \end{aligned}$$

Here, we indexed the subsets $T \subseteq V$ by their sizes.

For any node set T of size k , we can upper-bound the number of hyperedges contained within T . By using the fact that for any set T with $|T| = k$, there are at least k hyperedges incident to nodes in T , we have:

$$\begin{aligned} \|P^t - \pi\|_{TV} &\leq \sum_{k=1}^n \left(1 - \frac{k}{m}\right)^t \binom{n}{k} 2^{n-k} \\ &\leq n^2 \left(1 - \frac{1}{m}\right)^t \\ &\leq e^{-c} \end{aligned}$$

since for $f(k) = \left(1 - \frac{k}{m}\right)^t \binom{n}{k} 2^{n-k}$, we have $f(k) \geq f(k+1)$ for $k \geq 1$ and $t > 2m \log n + cn$. The theorem is proved. \square

C Proof of Theorem 7

Proof. Because the random walk on H^* is given by memoryless strategies, we can write its transition matrix as $\beta P_1 + (1 - \beta)P_2$ for memoryless P_1 . Theorem 3 allows us to analyze the eigenvalues of P_1 ; we can use results from [5] to see that P_2 has all eigenvalues between 0 and 1. Thus, if λ_T is an eigenvalue of P_1 , then there is a corresponding eigenvalue of $P = \beta P_1 + (1 - \beta)P_2$ satisfying:

$$\lambda \leq \beta \lambda_T + (1 - \beta).$$

Using Theorem 4, we have:

$$\begin{aligned}
|\mathbf{E}_t[A] - \pi(A)| &\leq \max_y \left| \sum_{x \in A} P^t(y, x) - \pi(x) \right| \\
&= \|P^t - \pi\|_{TV} \\
&\leq \sum_{T \subseteq V} \lambda_T^t (r-1)^{n-|T|} \\
&\leq \sum_{k=1}^n \left(1 - \frac{\beta k}{m}\right)^t \binom{n}{k} (r-1)^{n-k} \\
&\leq n^2 \left(1 - \frac{\beta}{m}\right)^t \\
&\leq e^{-c}
\end{aligned}$$

if

$$t \geq \frac{2m \log n + cn}{\beta}. \tag{2}$$

□