

Adversarial Scheduling Analysis of Game-Theoretic Models of Norm Diffusion

Gabriel Istrate^{1*}, Madhav V. Marathe², and S.S. Ravi³

¹ e-Austria Institute, V.Pârvan 4, cam. 045B, Timișoara RO-300223, Romania

² Network Dynamics and Simulation Science Laboratory, Virginia Tech, 1880 Pratt Drive Building XV, Blacksburg, VA 24061. Email: mmarathe@vbi.vt.edu

³ Computer Science Dept., S.U.N.Y. at Albany, Albany, NY 12222, U.S.A.
Email: ravi@cs.albany.edu

Abstract. In [IMR01] we advocated the investigation of robustness of results in the theory of learning in games under adversarial scheduling models. We provide evidence that such an analysis is feasible and can lead to nontrivial results by investigating, in an adversarial scheduling setting, Peyton Young’s model of diffusion of norms [You98]. In particular, our main result incorporates *contagion* into Peyton Young’s model.

Keywords: evolutionary games, adversarial scheduling, discrete Markov chains.

1 Introduction

Game-theoretic equilibria are *steady-state properties*; that is, given that all the players’ actions correspond to an equilibrium point it would be irrational for any of them to deviate from this behavior when the others stick to their strategy. The fundamental problem facing this type of concept is that it does not predict *how players arrive at this equilibrium in the first place*, or how they “choose” one such equilibrium, if several such points exist. The theory of equilibrium selection of Harsányi and Selten assumes some form of prior coordination between players, in the form of a *tracing procedure*. This strong prerequisite is often unrealistic.

The theory of *learning in games* [FL99] attempts to explain the emergence of equilibria as the result of an evolutionary “learning” process. Models of this type assume one (or several) populations of *agents*, that interact by playing a certain game, and updating their behavior based on the outcome of this interaction.

Results in evolutionary game theory are important not necessarily as realistic models of strategic interaction, but since they provide possible explanations for experimentally observed features of real-world social dynamics. For instance, the fundamental insight behind the concept of *stochastically stable strategies* is that continuous “noise” (or small deviations from rationality) can provide a solution to the equilibrium selection problem in game theory. Similar issues apply when mathematical modeling is replaced with computer experiments, in the area

* corresponding author. Email: gabrielistrate@acm.org

of *agent-based social simulation* [GT05]. Epstein [Eps07] (see also [AE96]) has advocated a generative approach to social science: in order to better understand a given phenomenon one should be able to generate it via simulations.

Given that such mathematical models or simulations are emerging as tools for policy-making (see e.g. [NBB99,ECC⁺04]), how can we be sure that the conclusions that we derive from the output of the simulation do not crucially depend on the particular assumptions and features we embed in it ? Part of the answer is that these results have to display “robustness” with respect to the various idealizations inherent in the model, be it mathematical or computational.

In this paper we are only concerned with one such issue: *scheduling*, i.e.the order in which agents get to update their strategies. Two alternatives are most popular, both in the mathematical and the computer simulation literature: in the *synchronous mode* (*every* player updates at every step. A popular alternative is *uniform matching*. Models of the latter type assume an underlying (hyper)graph topology (describing the sets of players allowed to simultaneously update in one step as a result of game playing) and choose a (hyper)edge uniformly at random from the available ones. Employing a uniform matching model in multiagent models of social systems is unrealistic for it assumes *perfect* and *global* randomness; it is not clear whether this assumption is warranted in the “real life” situations that the theory is supposed to model. A random matching model does not take into account *locality* and cannot, therefore, adequately model “contagion” effects (i.e. players becoming activated as a result of some of their neighbors doing so). On the other hand, social systems are *inherently distributed*, and it is not clear whether the assumption of *global* randomness is reasonable in a simulation setting.

We investigate in an adversarial setting Peyton Young’s model of evolution of norms [You98] (see also [You03]). The dynamics, although very simple, models an important aspect of social networks, the emergence of *conventions*, and has been proposed as an evolutionary justification for the emergence of certain rules in the pragmatics of natural language [Roo04]. Our results can be summarized as follows: results on selection of strict-dominant equilibria under random noise extend (Theorem 1) to a class of nonadaptive schedulers. However, such an extension fails (Theorem 2) for adaptive schedulers, even those with fairness properties similar to those of a random scheduler. Our main result (Theorem 3) extends the convergence to the strictly-dominant equilibrium to a class of “nonmalicious” adaptive schedulers that models *contagion* and has a certain *reversibility* property (the class of such schedulers includes the random scheduler as a special case). However for this class of schedulers the *convergence time* is *not* necessarily the one from the case of random scheduling.

Besides the relevance of our results to evolutionary game-theory, we hope that the concepts and techniques relevant to this paper can be fruitfully exploited in the theory of *rapidly mixing Markov chains*, of great interest in Theoretical Computer Science.

2 Preliminaries

A general class of models for which adversarial analysis can be naturally considered is that of *population games* [Blu01]. Systems of interest in this class consist of a number of *agents*, defined as the vertices of a hypergraph $H = (V, E)$. One edge of this hypergraph represents a particular choice of all players who can play (one or more simultaneous instances of) a game G that defines the dynamics. Each player has a *state* (generally a mixed strategy of G) chosen from a certain set S . The global state of the system is an element of $\bar{S} = S^V$. The dynamics proceeds by choosing one edge e of H (according to a scheduling mechanism that is specified by the scheduler), letting the agents in e play the game, and updating their states as a result of game playing.

2.1 Schedulers

Denote by X^* the set of finite words over alphabet X .

Definition 1. A (probabilistic) scheduler assigns a probability distribution $p_{w,s}$ on E to each pair (w, s, s_0) consisting of initial prefixes $w \in E^*$, $s \in \bar{S}^*$ with $|w| = |s|$ and starting state $s_0 \in \bar{S}$. The next element $e \in E$ to be scheduled, given prefixes w, s and initial state s_0 , is sampled from E according to p_{w,s,s_0} .

A non-adaptive probabilistic scheduler is specified by (a) a collection (multiset) $\Sigma = \{\mathcal{D}_1, \dots, \mathcal{D}_m\}$ of probability distributions on the set E such that every $x \in E$ belongs to the support of some distribution \mathcal{D}_i and (b) a fixed permutation π of Σ . The scheduler proceeds by (possibly concurrently) scheduling elements of E sampled from a distribution from Σ chosen according to (consecutive repetitions of) permutation π . For $C > 0$, a non-adaptive probabilistic scheduler is C individually-fair if for every $x \in E$, the probability that x is scheduled during one round of π is at least $C/|E|$.

One can define, for any given triple (w, s, s_0) , where $w \in E^*$, $s \in \bar{S}^*$ and $s_0 \in \bar{S}$, a probability π_{w,s,s_0} , the probability that, starting from state s_0 the scheduler uses w as the initial prefix of its schedule and evolves its global state according to string s . Let Ω denote the resulting probability space. We divide each trajectory of a probabilistic scheduler into *rounds*: the first round is the smallest initial segment that schedules each element of E at least once, the second round is the smallest segment starting at the end of the first round that schedules each element at least once, and so on.

Definition 2. If $f(\cdot)$ is a function on integers, we say that a family of probabilistic schedulers, indexed by n , the cardinality of the set E , is $O(f(n))$ -fair w.h.p. if there exists a monotonically decreasing function $g : (0, \infty) \rightarrow (0, 1)$, with $\lim_{\epsilon \rightarrow \infty} g(\epsilon) = 0$ such that for every state $s \in \bar{S}$, denoting by l_i the random variable measuring the length of the i 'th round, we have $\lim_{n \rightarrow \infty} \text{Prob}[l_i > \epsilon \cdot f(n)] < g(\epsilon)$.

Random scheduling can be specified by a non-adaptive probabilistic scheduler whose set Σ consists of just one distribution, namely the uniform distribution on E . This scheduler is 1-individually fair and, by the well-known Coupon Collector Lemma it is also $O(n \log(n))$ -fair w.h.p.

2.2 Peyton Young’s model of norm diffusion

The setup of this model is the following: agents located at the vertices of a graph G interact by playing a two-person symmetric game with payoff matrix $M = (m_{i,j})_{i,j \in \{\mathbf{A}, \mathbf{B}\}}$ displayed in Figure 1. It is assumed that strategy \mathbf{A} is a so called *strict risk-dominant equilibrium*. Without going to further details, in the context of this paper this assumption simply amounts to assuming that $a - d > b - c > 0$. Each undirected edge $\{i, j\}$ has a positive weight $w_{ij} = w_{ji}$ that measures its “importance”. When scheduled, agents play (using the same strategy, identified as *the agent’s state*) against each of their neighbors. If agent i is the one to update, \bar{x} is the joint profile of agents’ strategies, and $z \in \{\mathbf{A}, \mathbf{B}\}$ is the candidate new state, $p^\beta(x_i \rightarrow z | \bar{x}) \sim e^{\beta \cdot \nu_i(z, \bar{x}_{-i})}$, where $\nu_i(z, \bar{x}_{-i})$, the payoff of the i ’th agent should he play strategy z while the others’ profile remains the same is given by $\nu_i(z, \bar{x}_{-i}) = \sum_{(i,j) \in E} w_{ij} m_{z,x_j}$. Under random scheduling, the process we defined is a variant of the best-response dynamics. This latter process (viewed as a Markov chain) is not ergodic. Indeed, the since in game G it is always better to play the same strategy as your partner, the dynamics has at least two fixed points, states “all \mathbf{A} ” and “all \mathbf{B} ”.

strategies	\mathbf{A}	\mathbf{B}
\mathbf{A}	a,a	c,d
\mathbf{B}	d,c	b,b

Fig. 1. Payoff matrix

An important property of Peyton Young’s dynamics is that it corresponds to a *potential game*: there exists a function $\rho : \bar{V} \rightarrow \mathbf{R}$ such that, for any player i , any possible actions a_1, a_2 of player i , and any action profile \bar{a} of the other players, $u_i(a_1, \bar{a}) - u_i(a_2, \bar{a}) = \rho(a_1, \bar{a}) - \rho(a_2, \bar{a})$ (where u_i is the utility function of player i). In other words changes in utility as a result of strategy update correspond to changes in a global potential function. An explicit potential is given by $\rho^*(x) = \sum_{(h,k) \in E} w_{h,k} m_{x_h, x_k}$.

2.3 Stochastic stability

A fundamental concept we are dealing with is that of a *stochastically stable state* for dynamics described by a Markov chain.

Definition 3. Consider a Markov process P^0 defined on a finite state space Ω . For each $\epsilon > 0$, define a Markov process P^ϵ on Ω . P^ϵ is a regular perturbed Markov process if all of the following conditions hold.

- P^ϵ is irreducible for every $\epsilon > 0$.
- For every $x, y \in \Omega$, $\lim_{\epsilon \rightarrow 0} P_{xy}^\epsilon = P_{xy}^0$.
- If $P_{xy}^0 > 0$ then there exists $r(m) > 0$, the resistance of transition $m = (x \rightarrow y)$, such that as $\epsilon \rightarrow 0$, $P_{xy}^\epsilon = \Theta(\epsilon^{r(m)})$.

Let μ^ϵ be the (unique) stationary distribution of P^ϵ . A state s is stochastically stable if $\lim_{\epsilon \rightarrow 0} \mu^\epsilon(s) > 0$.

Peyton Young’s model of diffusion of norms can be easily recast into the framework of Definition 3, by defining $\epsilon = \exp(-\beta)$.

Definition 4. A tree rooted at node j is a set T of edges such that for any state $w \neq j$ there exists a unique (directed) path from w to j . The resistance of a rooted tree T is the sum of resistances of all edges in T .

The following characterization of stochastically stable states is presented as Lemma 3.2 in the Appendix of [You98]:

Proposition 1. Let P^ϵ be a regular perturbed Markov process, and for each $\epsilon > 0$ let μ^ϵ be the unique stationary distribution of P^ϵ . Then $\lim_{\epsilon \rightarrow 0} \mu^\epsilon = \mu^0$ exists, and μ^0 is a stationary distribution of P^0 . The stochastically stable states are precisely those states z such that there exists a tree rooted at z of minimal resistance (among all rooted trees).

Definition 5. Given a graph G , a nonempty subset S of vertices and a real number $0 \leq r \leq 1/2$ we say that S is r -close-knit if $\forall S' \subseteq S, S' \neq \emptyset$, $\frac{e(S', S)}{\sum_{i \in S'} \deg(i)} \geq r$, where $e(S', S)$ is the number of edges with one endpoint in S' and the other in S , and $\deg(i)$ is the degree of vertex i . A graph G is (r, k) -close-knit if every vertex is part of a r -close-knit set S , with $|S| = k$.

Definition 6. Given $p \in [0, 1]$, the p -inertia of the process is the maximum, over all states $x_0 \in S$, of $W(\beta, p, x_0)$, the expected waiting time until at least $1 - p$ of the population is playing action A conditional on starting in state x_0 .

The model in [You03, You98] assumes independent individual updates, arriving at random times governed (for each agent) by a Poisson arrival process with rate one. Since we are, however, interested in adversarial models that do not have an easy description in continuous time we will assume that the process proceeds in discrete steps. At each such step a random node is scheduled. It is a simple exercise to translate the result in [You03, You98] to an equivalent one for global, discrete-time scheduling. The conclusions of this translation are:

- The stationary distribution of the process is the Gibbs distribution, $\mu_\beta(x) = \frac{e^{\beta \rho(x)}}{\sum_z e^{\beta \rho(x)}}$, where ρ is the potential function of the dynamics.
- “All \mathbf{A} ” is the unique stochastically-stable state of the dynamics.
- Let $r^* = \frac{b-c}{a-d+b-c}$, and let $r > r^*$, $k > 0$. On a family of (r, k) -close-knit graphs the convergence time is $O(n)$.

3 Results

First we show that Peyton Young’s results extend to non-adaptive schedulers:

Theorem 1. *The following hold: (i) For all non-adaptive schedulers, the state “all **A**” is the unique stochastically stable state of the system. (ii) Let \mathcal{G} be a class of graphs that are (r, k) -close-knit for some fixed $r > r^*$. Let $f = f(n)$ be a class of non-adaptive $\Theta(1)$ individually fair schedulers. Given any $p \in (0, 1)$ there exists a β_p such that for all $\beta > \beta_p$ there exists a constant C such that the p -inertia of the process (under scheduling given by f) is at most $C \cdot m \cdot n$, where $m = m(n)$ is the number of rounds of f and n is the number of vertices of the underlying graph.*

Adaptive schedulers on the other hand, even those of fairness no higher than that of the random scheduler, can preclude the system from ever entering a state where a proportion higher than r of agents plays the risk-dominant strategy:

Theorem 2. *For every $0 < r < 1$ there exists an adaptive scheduler which is $O(n \log(n))$ -fair w.h.p. (where the constant hidden in the “ O ” notation depends on r) that can forever prevent the system, started on the “all **B**’s” configuration, from ever having more than a fraction of r of the agents playing **A**.*

3.1 Diffusion of norms by contagion

Adaptive schedulers can display two very different notions of adaptiveness:

- (i) The next node depends only on the set of previously scheduled nodes, or
- (ii) It crucially depends on the *states* of the system so far.

The adaptive schedulers in Theorem 2 was crucially using the second, stronger, kind of adaptiveness. In the sequel we study a model that displays adaptiveness of type (1) but not of type (2). The model is specified as follows: To each node v we associate a probability distribution D_v on the vertices of G . We then choose the next scheduled node according to the following process. If t_i is the node scheduled at stage i , we chose t_{i+1} , the next scheduled node, by sampling from D_{t_i} . In other words, the scheduled node performs a (non-uniform) random walk on the vertices of graph G . To exclude technical problems such as the periodicity of this random walk, we assume that it is always the case that $v \in \text{supp}(D_v)$. Also, let H be the directed graph with edges defined as follows: $(x, y) \in E[H] \iff (y \in \text{supp}(D_x))$. This dynamics generalizes both the class of non-adaptive schedulers from previous result and the random scheduler (for the case when H is the complete graph). In the context of van Roooy’s evolutionary analysis of signalling games in natural language [Roo04], it functions as a simplified model for an essential aspect of emergence of linguistic conventions: transmission via *contagion*.

It is easy to see that the dynamics can be described by an aperiodic Markov chain M on the set on $V^{\{\mathbf{A}, \mathbf{B}\}} \times V$, where a state (\bar{w}, x) is described as follows:

- \bar{w} is the set of strategies chosen by the agents.
- x is the label of the last agent that was given the chance to update its state.

If the directed graph H is strongly connected then the Markov chain M is irreducible, hence it has a stationary distribution Π . We will, therefore, limit ourselves in the sequel to settings with strongly connected H . We will, further, assume that the dynamics is *weakly reversible*, i.e. $(x \in \text{supp}(D_y))$ if and only if $(y \in \text{supp}(D_x))$. This, of course, means that the graph H is undirected. Note that since we do not constrain otherwise the transition probabilities of distributions D_i , the stationary distribution Π of the Markov chain does *not*, in general, decompose as a product of component distributions. That is, one cannot generally write $\Pi(w, x)$ as $\Pi(w, x) = \pi(w) \cdot \rho(x)$, for some distributions π, ρ .

Theorem 3. *The set $Q = \{(w, x) | w = V^A\}$ is the set of stochastically stable states for the diffusion of norms by contagion.*

Proof. The states in Q are obviously reachable from one another by zero-resistance moves, so it is enough to consider one state $y \in Q$ and prove that it is stochastically stable. To do so, by Proposition 1, all we need to do is show that y is the root of a tree of minimal resistance. Indeed, consider another state $x \in Q$ and let T be a minimum potential tree rooted at x .

Claim. There exists a tree \bar{T} rooted at y having potential less or equal to the potential of the tree T , strictly smaller in case x is not a state having all its first-component labels equal to A .

Let $\pi_{y,x} = (x_0, i_0) \rightarrow (x_1, i_1) \rightarrow \dots \rightarrow (x_k, i_k) \rightarrow (x_{k+1}, i_{k+1}) \rightarrow \dots \rightarrow (x_r, i_r)$ be the path from y to x in T (that is $(x_0, i_0) = y, (x_r, i_r) = x$).

We will define \bar{T} by viewing the set of edges of T as partitioned into subsets of edges corresponding to paths as follows (see Figure 2 (a)):

- The set of edges of path $\pi_{y,x}$.
- The set of edges of the subtree rooted at y .
- Edges of tree components (perhaps consisting of a single node) rooted at a node of $\pi_{y,x}$, other than y (but possibly being x).

To obtain \bar{T} we will transform each tree (path) in the above decomposition of T into one that will be added to \bar{T} . The transformation goes as follows:

- Instead of path $\pi_{y,x}$ we add path $\Pi_{x,y}$ from x to y defined by: $\Pi_{x,y} = (x_r, i_r) \rightarrow (x_{r-1}, i_{r-1}) \rightarrow (x_{r-2}, i_{r-2}) \rightarrow \dots \rightarrow (x_0, i_1) \rightarrow (x_0, i_0) = y$.
- Rooted trees of type (2) are included into tree \bar{T} as well.
- The transformation is more complicated for the third type of edges, and we explain it in detail. Let W_k be a tree component of T , connected to path $\pi_{y,x}$ at connection point (x_k, i_k) .

Case 1: $x_k = x_{k-1}$. Then the point $(x_k, i_k) = (x_{k-1}, i_k)$ belongs to path $\Pi_{x,y}$ as well, so one can just add the rooted tree W_k to \bar{T} as well.

Case 2: $x_k \neq x_{k-1}$ and the move $(x_{k-1}, i_{k-1}) \rightarrow (x_k, i_k)$ has positive resistance. In this case, since in configuration x_{k-1} and scheduled node i_k we




Fig. 2. (a). Decomposition of edges of tree T (b). Resistance of edges on a path between two nodes X and Y .

have a choice of either moving to x_k or staying in x_{k-1} , it follows that the move $(x_k, i_k) \rightarrow (x_{k-1}, i_k)$ has zero resistance.

Therefore we can add to \bar{T} the tree $\bar{W}_k = W_k \cup \{(x_k, i_k) \rightarrow (x_{k-1}, i_k)\}$. The tree has the same resistance as the one of tree W_k .

Case 3: $x_{k-1} \neq x_k$ and the move $(x_{k-1}, i_{k-1}) \rightarrow (x_k, i_k)$ has zero resistance. Let j be the smallest integer such that either $x_{k+j+1} = x_{k+j}$ or $x_{k+j+1} \neq x_{k+j}$ and the move $(x_{k+j}, i_{k+j}) \rightarrow (x_{k+j+1}, i_{k+j+1})$ has positive resistance. In this case, one can first replace W_k by $W_k \cup \{(x_k, i_k) \rightarrow (x_{k+1}, i_{k+1}), (x_{k+1}, i_{k+1}) \rightarrow \dots \rightarrow (x_{k+j}, i_{k+j})\}$ without increasing its total resistance. Then we apply one of the techniques from Case 1 or Case 2.

Case 4: $x_{k-1} \neq x_k$, the move $(x_{k-1}, i_{k-1}) \rightarrow (x_k, i_k)$ has zero resistance, and all moves on $\pi_{y,x}$, from x_k up to x have zero resistance. Then define $\bar{W}_k = W_k \cup (x_k, i_k) \rightarrow (x_{k+1}, i_{k+1}) \rightarrow \dots \rightarrow x$.

It is easy to see that no two sets W_k intersect on an edge having positive resistance. The union of the paths of all the sets is a directed associated graph W rooted at y , that contains a rooted tree \bar{T} of potential no larger than the potential of W . Since transformations in cases (i),(iii) do not increase tree resistance, to compare the potentials of T and W it is enough to compare the resistances of paths $\pi_{y,x}$ and $\Pi_{x,y}$.

We come now to a fundamental property of the game G : since it is a potential game, the resistance $r(m)$ of a move $m = (a_1, j_1) \rightarrow (a_2, j_2)$ only depends on the values of the potential function at three points: a_1, a_2 and a_3 , where a_3 is the state obtained by assigning node j_2 the value not assigned by move to a_2 . Specifically, $r(m) > 0$ if either $\rho^*(a_2) < \rho^*(a_1)$, in which case $r(m) = \rho^*(a_1) - \rho^*(a_2)$, or $a_2 = a_1$ and $\rho^*(a_3) > \rho^*(a_1)$, in which case $r(m) = \rho^*(a_3) - \rho^*(a_1)$. In other words, the resistance of a move is positive in the following two cases: (1) The move leads to a decrease of the value of the potential function. In this case the resistance is equal to the difference of potentials. (2) The move corresponds to keeping the current state (thus not modifying the value of the potential function), but the alternate move would have increased the potential. In this case the resistance is equal to the value of this increase.

Let us now compare the resistances of paths $\pi_{y,x}$ and $\Pi_{x,y}$. First, the two paths contain no edges of infinite resistance, since they correspond to possible moves under Markov chain dynamics P^ϵ . If we discount second components, the two paths correspond to a single sequence of states Z connecting x_0 to x_r , more precisely to *traversing Z in opposite directions*. (The last move in $\Pi_{x,y}$ has zero resistance and can thus be discounted). Resistant moves of type (2) are taken into account by both traversals, and contribute the same resistance value to both paths. So, to compare the resistances of the two paths it is enough to compare resistance of moves of type (1). Moves of type (1) of positive resistance are those that lead to a decrease in the potential function. Decreasing potential in one direction corresponds to increasing it in the other (therefore such moves have zero resistance in the opposite direction).

An illustration of the two types of moves is given in Figure 2 (b), where the path between X and Y goes through four other nodes, labeled 1 to 4. The relative height of each node corresponds to the value of the potential function at that node. Nodes 2 and 3 have equal potential, so the transition between 2 and 3 contributes an equal amount to the resistance of paths in both directions (which may be positive or not). Other than that only transitions of positive resistance are pictured.

The conclusion of this argument is that $r(\pi_{y,x}) - r(\Pi_{x,y}) = \rho^*(x) - \rho^*(y) \geq 0$, and $r(\pi_{y,x}) - r(\Pi_{x,y}) > 0$ unless x is an “all A ” state.

3.2 The inertia of diffusion of norms with contagion

Theorem 3 shows that random scheduling is not essential in ensuring that stochastically stable states in Peyton Young’s model correspond to all players playing A : the same result holds in the model with contagion. On the other hand, the result on the p -inertia of the process on families of close-knit graphs is *not* robust to such an extension. Indeed, consider the *line graph* L_{2n+1} on $2n+1$ nodes labelled $-n, \dots, -1, 0, 1 \dots n$. Consider a random walk model such that: (a) the origin of the random walk is node 0, and (b) the walk goes left, goes right or stays in place, each with probability $1/3$. It is a well-known property of the random walk that it takes $\Omega(n^2)$ time to reach nodes at distance $\Omega(n)$ from the origin. Therefore, the p -inertia of this random walk dynamics is $\Omega(n^2)$ even though for every $r > 0$ there exists a constant k such that the family $\{L_{2n+1}\}$ is (r, k) -close-knit for large enough n .

In the journal version of the paper we will present an upper bound on the p -inertia for the diffusion of norms with contagion based on concepts similar to the *blanket time* of a random walk [WZ96].

4 Conclusions and Acknowledgments

Our results have made the original statement by Peyton Young more robust, and have highlighted the (lack of) importance of various properties of the random scheduler in the results from [You98]: the *reversibility* of the random scheduler,

as well as its inability to use the global system state are important in an adversarial setting, while its fairness properties are not crucial for convergence, only influencing convergence time. Also, the fact that the stationary distribution of the perturbed process is the Gibbs distribution (true for the random scheduler) does not necessarily extend to the adversarial setting.

This work has been supported by the Romanian CNCSIS under a PN-II “Idei” Grant, by the U.S. Department of Energy under contract W-705-ENG-36 and by NSF Grant CCR-97-34936.

References

- [AE96] R. Axtell and J. Epstein. *Growing Artificial Societies: Social Science from the Bottom Up*. The MIT Press, 1996.
- [Blu01] L. Blume. Population games. In S. Durlauf and H. Peyton Young, editors, *Social Dynamics: Economic Learning and Social Evolution*. MIT Press, 2001.
- [ECC⁺04] J.M. Epstein, D. Cummings, S. Chakravarty, R. Singa, and D. Burke. *Toward a Containment Strategy for Smallpox Bioterror. An Individual-Based Computational Approach*. Brookings Institution Press, 2004.
- [Eps07] J. Epstein. *Generative Social Science: Studies in Agent-based Computational Modeling*. Princeton University Press, 2007.
- [FL99] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. M.I.T. Press, 1999.
- [GT05] N. Gilbert and K. Troizch. *Simulation for social scientists (second edition)*. Open University Press, 2005.
- [IMR01] G. Istrate, M.V. Marathe, and S.S. Ravi. Adversarial models in evolutionary game dynamics. In *Proceedings of the 13th ACM-SIAM Symposium on Discrete Algorithms (SODA'01)*, 2001. (journal version in preparation).
- [NBB99] K. Nagel, R. Beckmann, and C. Barrett. TRANSIMS for transportation planning. In Y. Bar-Yam and A. Minai, editors, *Proceedings of the Second International Conference on Complex Systems*. Westview Press, 1999.
- [Roo04] R. Van Rooy. Signalling games select Horn strategies. *Linguistics and Philosophy*, 27:423–497, 2004.
- [WZ96] P. Winkler and D. Zuckerman. Multiple cover time. *Random Structures and Algorithms*, 9(4):403–411, 1996.
- [You93] H.P. Young. The evolution of conventions. *Econometrica*, 61(1):57–84, 1993.
- [You98] H.P. Young. *Individual Strategy and Social Structure: an Evolutionary Theory of Institutions*. Princeton University Press, 1998.
- [You03] H.P. Young. The diffusion of innovations in social networks. In L. Blume and S. Durlauf, editors, *The Economy as a Complex System III*. Oxford University Press, 2003.

Appendix

4.1 Proof of Theorem 1

Proof. (i) Let $i \in V$ and let M_i be the restriction of the given dynamics corresponding to the case when only one node, node i , is scheduled in all moves (otherwise the dynamics is similar to the original one).

It is easy to see that M_i is a non-ergodic Markov chain and that μ_β is a stationary distribution for the Markov chain M_i . This is so because for two configurations x, y that only differ in position i , the ratio of transition probabilities $p_{x,y}/p_{y,x}$ is equal to $\exp[\beta \cdot (\rho^*(x) - \rho^*(y))]$, which is precisely $\mu_\beta(x)/\mu_\beta(y)$.

Now consider the matrix D_k corresponding to the distribution with the same notation as in the periodic schedule. It is a convex combination of the matrices M_i , hence it will also have μ_β as a stationary distribution. We infer that the product of matrices D_k corresponding to the cyclic schedule also has μ_β as a stationary distribution.

But it is easy to see that the Markov chain corresponding to one round of the cyclic schedule is irreducible (since one can navigate between any two states in at most $|V|$ rounds, by flipping the differing bits and keeping the bits that coincide fixed) and aperiodic (since the probability of remaining in a given state is positive). Therefore, it must have a unique stationary distribution, which is necessarily μ_β .

- (ii) Consider a vertex $v \in V$ and a r -close-knit set of size k containing v , denoted S_v . Consider $\Gamma_{v,\beta}$ the version of the process where each vertex in S_v updates just as before, but each vertex in $V \setminus S_v$ always chooses state B when scheduled.

This restricted dynamics on V still corresponds to a potential game, specified by potential function

$$\rho^*(x) = \sum_{(i,j) \in E} \rho(x_i, x_j), \text{ for } x \in \{A, B\}^{S_v} B^{V \setminus S_v},$$

and with the Gibbs distribution $\mu^\beta(x) = \frac{e^{\beta \cdot \rho^*(x)}}{\sum_z e^{\beta \cdot \rho^*(z)}}$ as its stationary distribution. Again, just as in [You93] (since the precise scheduling order does *not* play a role in this result) the condition that G is (r, k) -close-knit implies that the state A_S defined as “all **A**” on S_v and “all **B**” on $V \setminus S_v$ is the state with the highest potential among the possible states of the system.

One additional complication of the dynamics $\Gamma_{v,\beta}$ is that it often schedules (unnecessarily) nodes outside S_v , that do not change. Consider $\Xi_{v,\beta}$ that is the version of $\Gamma_{v,\beta}$ that “only schedules nodes in S_v ” (i.e. it ignores moves of $\Gamma_{v,\beta}$ that schedule nodes outside of S_v).

To describe this dynamics formally, view each distribution D_i as a set of symbols from the alphabet V . Then the set of trajectories of the dynamics $\Gamma_{v,\beta}$ can be specified by the words of the regular language $L_\Gamma = (D_1 \cdot D_2 \cdot \dots \cdot D_m)^*$. Trajectories of $\Xi_{v,\beta}$ correspond to words in another regular language

L_{Ξ} , more precisely to the ones corresponding to deleting symbols in $V \setminus S_v$ from words in L_{Γ} . This regular language can be specified by the regular expression $((D_1 \cup \{\lambda\}) \cdot \dots \cdot (D_m \cup \{\lambda\}) \cap S_v^+)^*$. This expression yields a matrix of size $2^{|S_v|} \times 2^{|S_v|}$ for $\Xi_{v,\beta}$.

Claim. For every $\epsilon > 0$ there exists $\eta \in \mathbf{N}$ such that, for every $\eta' > \eta \in \mathbf{N}$ any initial state of $\Gamma_{v,\beta}$ and every state $T \in \{\mathbf{A}, \mathbf{B}\}^{S_v}$.

$$|Pr[\Gamma_{v,\beta} \text{ in state } T \mid |w| = \eta' \cdot m \cdot n] - \Pi(T)| \leq \epsilon.$$

Proof. Let $\epsilon > 0$. As $\Xi_{v,\beta}$ converges to its stationary distribution, there exists $\bar{k} > 0$ such that $\forall k' > \bar{k}$ and every initial state of $\Xi_{v,\beta}$

$$|Pr[\Xi_{v,\beta} \text{ in state } T \mid |w| = k'] - \tilde{\Pi}(T)| \leq \epsilon/2, \quad (1)$$

where $\tilde{\Pi}$ is the stationary distribution of dynamics Ξ . Of course, states with positive support in $\tilde{\Pi}$ have the same probability in Π , that is

$$\forall T \in \{\mathbf{A}, \mathbf{B}\}^{S_v}: \Pi(T) = \tilde{\Pi}(v).$$

Let Y be a random trajectory of length $\eta' \cdot m \cdot n$ in l_{Γ} and let $pr(Y)$ its projection onto L_{Ξ} .

Claim. There exists $\eta > 0$ such that $\forall \eta' > \eta$

$$\text{Prob}_{|Y|=\eta' \cdot m \cdot n} [|pr(Y)| < \bar{k}] \leq \frac{\epsilon}{2}.$$

Proof. The probability that any given distribution D_i whose support includes some element in S_v will schedule (in a given round) a node in this set is $\Omega(1/n)$, by the fairness condition. There is at least one such D_i among all the m distributions. Therefore, the expected length of $pr(Y)$ is $\Omega(k/n)$. A simple application of Markov's inequality gives the desired result. □

Now write

$$\begin{aligned} & Pr [\Gamma_{v,\beta} \text{ in state } T \mid |w| = \eta' \cdot m \cdot n] \\ &= \sum_j Pr[\Gamma_{v,\beta} \text{ in state } T \mid |w| = \eta' \cdot m \cdot n, |pr(w)| = j] \\ & \quad \cdot \text{Prob}[|pr(w)| = j] \end{aligned}$$

Therefore we have

$$\begin{aligned} & Pr [\Gamma_{v,\beta} \text{ in state } T \mid |w| = \eta' \cdot m \cdot n] - \Pi(T) \leq \\ & \leq \sum_j |Pr[\Gamma_{v,\beta} \text{ in state } T \mid |w| = \eta' \cdot m \cdot n, |pr(w)| = j] - \Pi(T)| \cdot \\ & \quad \cdot \text{Prob}[|pr(w)| = j] \end{aligned}$$

The first term in the product is an absolute difference between two probability values, and thus has absolute value at most one. Therefore, by Claim ii, if we neglect in the sum those terms with $j < k$ only changes the sum by at most $\epsilon/2$. On the other hand

$$\begin{aligned} Pr[\Gamma_{v,\beta} \text{ in state } T \mid |w| = \eta' \cdot m \cdot n, |pr(w)| = j] \\ = Pr[\Xi_{v,\beta} \text{ in state } T \mid |w| = k']. \end{aligned}$$

Now, using Equation (1), Claim ii follows. □

From now on the proof mirrors rather closely the one for the case of random scheduling (presented in the Appendix to [You93]): first, because the stationary distribution of process Γ is the Gibbs distribution, there exists a finite value $\beta(\Gamma, S, p)$ such that $\mu_v(A_S) \geq 1 - p^2/2$ for all $\beta > \beta(\Gamma, S, p)$. There are only a finite number of nonisomorphic dynamical systems $\Gamma_{v,\beta}$ (where isomorphism of dynamical system is meant to be the isomorphism of the underlying graph topologies S_v and of the projection of schedulers onto S_v). In particular we can find $\beta(r, k, p)$ and $\eta(r, k, p)$ such that, for all graphs $G \in \mathcal{G}$ and all r -close-knit subsets S with k vertices, the following relation holds for *all* initial states:

$$\forall \beta \geq \beta(r, k, p), \forall \eta' \geq \eta, Pr[y^{\eta' \cdot m \cdot n} = A_S] \geq 1 - p^2. \quad (2)$$

where y^t is the state of the dynamical system on state S_v at time t . We can now derive that for every close-knit set S

$$\forall \beta \geq \beta(r, k, p), \forall \eta' \geq \eta, Pr[x^{\eta' \cdot m \cdot n} = A_S] \geq 1 - p^2. \quad (3)$$

where x_t is now the state of the process from the theorem. The argument is obtained via essentially the same coupling as the one from [You98], hence it is omitted from this writeup. Since every vertex i is contained in a (r, k) -close-knit set, it follows that

$$\forall \beta \geq \beta(r, k, p), \forall \eta' \geq \eta, Pr[x_i^{\eta' \cdot m \cdot n} = \mathbf{A}] \geq 1 - p^2.$$

Therefore the expected proportion of vertices playing action A at time $\eta' \cdot m \cdot n$ is at least $(1 - p^2)n$.

But this implies that

$$\forall t \geq \eta \cdot m \cdot n, Pr[\text{at least } (1 - p)n \text{ nodes have label } \mathbf{A}] \geq 1 - p.$$

Indeed, if this wasn't the case, then with probability at least p more than pn nodes at time t would have label B , a contradiction. Now, by an application of Markov's inequality, the expected time until at least $(1 - p)n$ nodes are labelled A is at most $\eta \cdot m \cdot n / (1 - p)$. Since this holds for all graphs G in \mathcal{G} , the p -inertia of the process is bounded as stated in the theorem. □

4.2 Proof of Theorem 2

Proof. Consider a scheduler working in rounds. In each round the scheduler is scheduling nodes according to a fixed permutation π , the same for all rounds. In each round the scheduler is scheduling each node at least once. For the first $\lceil rn \rceil + 1$ nodes the scheduler continues scheduling each of them (after the initial one) until the node switches to strategy **B**. The scheduler plays each remaining node *exactly* once.

It is easy to see that there exists a constant $\epsilon > 0$ (that may depend on β) such that, at each stage, each agent switches to strategy **B** with probability greater or equal to ϵ .

Therefore the probability that any given agent needs to be scheduled for more than $c \log(n)$ rounds before turning to **B** is $o(1/n)$ for large enough c . It follows that the given scheduler is $O(n \log(n))$ -fair w.h.p.

□