Dynamics in Network Interaction Games *

Martin Hoefer † Siddharth Suri ‡

November 17, 2011

Abstract

We study the convergence times of dynamics in games involving graphical relationships of players. Our model of local interaction games generalizes a variety of recently studied games in game theory and distributed computing. In a local interaction game each agent is a node embedded in a graph and plays the same 2-player game with each neighbor. He can choose his strategy only once and must apply his choice in each 2-player game he is involved in. This represents a fundamental model of decision making with local interaction and distributed control. Furthermore, we introduce a generalization called 2-type interaction games, in which one 2-player game is played on edges and possibly another game is played on non-edges. For the popular case with symmetric 2×2 games, we show that several dynamics converge to a pure Nash equilibrium in polynomial time. This includes arbitrary sequential better-response dynamics, as well as concurrent dynamics resulting from a distributed protocol that does not rely on global knowledge. We supplement these results with an experimental comparison of sequential and concurrent dynamics.

*An extended abstract of this work has appeared in DISC 2009 [26].

†Department of Computer Science, RWTH Aachen University, Germany. mhoefer@cs.rwth-aachen.de. Supported by a fellowship within the Postdoc-Program of the German Academic Exchange Service (DAAD), by DFG through UMIC Research Center at RWTH Aachen University, and by DFG grant Ho 3831/3-1.

‡Yahoo! Research, New York, USA. suri@yahoo-inc.com.

1 Introduction

In this paper we examine convergence of dynamics in a fundamental model for distributed decision making with local interactions. We introduce two game-theoretic models, one a generalization of the other, that combine strategic interaction with the notion of graph-based locality. This extends a variety of game-theoretic settings that have been studied intensively in the literature. In our model of a *local interaction game* there is a graph G along with a 2-player symmetric game, Γ. Players are the nodes, and the graph models the local interaction possibilities. In particular, Γ is played along each edge of G, and each player plays the same strategy against each of their neighbors. The payoff of a player is simply the sum of the payoffs earned from playing each neighbor. We also introduce a generalization of local interaction games called *2-type interaction games*. Intuitively, a 2-type interaction game is a graph where one 2-player symmetric game is played on the edges, and another 2-player symmetric game is played on the non-edges. Whereas local interaction games model the restricted interaction *possibilities* of players through the topology of the graph, 2-type interaction games also model different *types* of interactions that occur between players. This is a natural assumption when considering e.g. social networks, as they do not necessarily indicate restrictions of interactions, but rather show that there is a special relationship, which is likely to alter the incentives of the involved actors. Our model allows one to specify for example how one person treats a friend differently than a stranger. In addition, it is possible to study distributed graph clustering problems (such as correlation clustering [22]) within this framework.

Our main interest is how the set of players can quickly arrive at a stable set of decisions—a Nash equilibrium of the game—using distributed decision making policies. Our main result is that myopic sequential better-response dynamics converge in polynomial time to a Nash equilibrium. This results holds for 2-type interaction games based on arbitrary symmetric 2×2 games and arbitrary graphs, which encompasses the vast majority of cases considered previously in related work. More interestingly, we can show a similar result for a payoff-relative concurrent protocol without central coordination. This result holds for arbitrary local interaction games, and 2-type interaction games with polynomially bounded payoffs. While sequential better response has a natural and intuitive appeal, our concurrent policy is carefully designed. It exhibits a number of favorable properties, such as respecting player incentives and relying only on local information. Designing such policies that yield provably rapid convergence is a major concern in wireless networks and distributed control systems (see, e.g. [17, 34]), and our results contribute to this research agenda. In particular, our dynamics can be interpreted as efficient distributed algorithms to compute a Nash equilibrium, which stands in sharp contrast to other game-theoretic models of restricted (graphical) interaction [13, 30].

The comparison of convergence times for sequential and concurrent dynamics in local interaction games without a dominant strategy reveals that the lack of central control can result in concurrent dynamics being slower than sequential ones. This, however, is a worst-case result, and we indicate that in coordination games, where the incentive is to join neighbors in their strategy choice, concurrent dynamics resulting from our protocol can be significantly faster. This does not necessarily hold for anti-coordination games, where players try to pick different strategies than their neighbors. Here a simple adjustment of our concurrent dynamics to a fixed choice μ for the migration probability can yield better results. However, the choice of this value is delicate, as resulting dynamics might abruptly drown in oscillation. It remains an interesting open problem to find improved analytical bounds for expected convergence times in specific classes of local interaction games.

The rest of the paper is structured as follows. We revisit related work in Section 1.1 and define the model in Section 1.2. Sequential dynamics are treated in Section 2, concurrent dynamics in Section 3. In Section 4 we compare sequential and concurrent convergence times in simple local interaction games. Finally, Section 5 concludes the paper.

1.1 Related Work

This paper fits into a recent stream of works that study subclasses of local interaction games. For example, our model of local interaction games generalizes a game considered by Bramoullé [9], which concentrates on the subclass of symmetric 2×2 anti-coordination games on the edges and does not have any games on the non-edges. In the rest of this section we discuss models that are related to, but are not necessarily specializations of local interaction games. A special class of anti-coordination game derived from the MaxCut problem has been used in [16]. It was studied by Christodoulou et al. [12] in terms of convergence time to Nash equilibria and social welfare of states obtained after a polynomial number of best-response steps.

Variants of local interaction games with coordination games are central in the study of threshold phenomena, cascading dynamics, and information diffusion in networks [33]. Closest to our focus is a recent paper by Montanari and Saberi [37] who consider a special case of our model of local interaction games, in which the 2×2 games are symmetric coordination games. For these games, they study a class of noisy best-response dynamics called logit-response, heat bath, or Glauber dynamics. For potential games, in the long run, the time this process spends at a state scales proportional to the potential value and the noise level. For small noise levels the dynamics thus remain exclusively at global potential maximizers. For coordination games this is a state in which all players use the same strategy, and in such games it is also the state maximizing the social welfare (defined as the sum of player utilities). Studying properties of these dynamics in this class of games has been a popular topic in numerous works over the last two decades [6,11,14,40]. The results of [37] are complementary to ours in the sense that they consider the hitting time of a *global* potential maximizer in a significantly more restrictive model. They show that convergence times increase from polynomial to exponential time when the graph becomes more well-connected. This contrasts our polynomial time bounds for all graphs and arbitrary symmetric 2×2 games when only convergence to *local* potential maximizers is required. Note that in our more general model, a potential maximizer can be arbitrarily bad in terms of social welfare (e.g., for 2×2 prisoner's dilemma games it is even a *minimizer* of social welfare). Naturally, in terms of stability, a potential maximizer remains a desirable state, however, the results of [37] indicate that quick convergence by better-response protocols to such states does not hold. Furthermore, in our more general scenario, even computing a potential maximizer is NP-complete, e.g., in the MaxCut game. This motivates us to derive protocols that guarantee convergence to a slightly less stringent notion of stability, i.e., to pure Nash equilibria.

Our concurrent dynamics are closely related to recent work on protocols for concurrent strategy updating in potential games for distributed control in networks [34,35], some of which are inspired by evolutionary game theory [1,17,18]. In addition, there is a large body of related work on strategic learning [21,41], various forms of dynamics such as calibrated [19] or regret learning [5,25,42] or best response/ficticious play [2,12,16,20,36], and a variety of equilibrium concepts such as correlated Nash [3] or sink equilibria [15,23]. In a related work, Kearns and Tan [32] design a voting protocol with polynomial time convergence in a similar 2-strategy coordination scenario. In contrast to our work they also require convergence to a state where all players play the same strategy. In the graphical model of evolutionary game theory introduced by Kearns and Suri [31] all players play a 2-player symmetric game with a randomly chosen neighbor. The authors characterize evolutionary equilibria in terms of the graph structure. However, they give no notion of dynamics that converge to equilibrium.

While in our model the graph is fixed and specified in advance, there are several works on games with network formation. In particular, 2×2 anti-coordination games on endogenous graphs were studied in [10]. More work has been done on network formation and 2×2 coordination games. In these games, players strategically decide both which links to build and which strategy to play in the games on the links. These games are classes of local interaction games with network creation, i.e., they allow only connected players to interact. There has been no focus on duration of dynamics, social welfare, and computation of Nash equilibria and optimal states. Instead, properties of the network structure and payoff properties in Nash equilibria were analyzed [4], or stochastically stable states were characterized [24,27].

1.2 Model and Notation

We begin by giving the formal definition of a 2-type interaction game.

Definition 1. *A* 2-type interaction game *is a graph $G = (V, E)$ together with two, possibly different, 2-player symmetric games Γ^c and Γ^d, where the set of strategies is the same in both games.*

Intuitively, on each edge $e \in E$ connected players play a 2-player symmetric game Γ^c. In addition, for each non-edge the pair of disconnected players play a possibly different symmetric 2×2 game, Γ^d. Each player plays the same strategy in all games he is participating. In this work, we restrict Γ^c and Γ^d to be 2×2 symmetric games with strategies 1 and 2, and payoffs for Γ^c and Γ^d are denoted as shown in Figure 1.

Next we introduce some notation that will allow us to define the utility function for each player in a 2-type interaction game. We will let $\Gamma_p = \{\Gamma^c, \Gamma^d\}$. We denote by $n = |V|$ the number of players, $m = |E|$

Γ^c	1	2
1	a,a	d,c
2	c,d	b,b

Γ^d	1	2
1	e,e	h,g
2	g,h	f,f

Figure 1: Payoffs in the Games.

the number of edges, and $\deg(v)$ the degree of player v. Let $S = \{1,2\}^n$ be the set of states of the game and $s = (s_v)_{v \in V} \in S$ a state, where $s_v \in \{1,2\}$ is the strategy of player v. For a state s the set of players playing strategy 1 is denoted $V_1(s)$, and their number is denoted $|s|_1$. For a player v the number $|s|_1^v$ denotes the number of her neighbors playing 1. V_2, $|s|_2$, $|s|_2^v$ are defined similarly for strategy 2. The size of the cut of a state s, which is the number of edges connecting players that play different strategies, is denoted by $m_{12}(s)$. A player v has utility for strategy 1

$$\begin{aligned} \text{util}_v(1, s_{-v}) \quad = \quad & [\mathtt{a} \cdot |s|_1^v & + \quad \mathtt{d} \cdot |s|_2^v] \\ + \quad & [\mathtt{e} \cdot (|s|_1 - 1 - |s|_1^v) & + \quad \mathtt{h} \cdot (|s|_2 - |s|_2^v)] \ , \end{aligned}$$

while for strategy 2 he has utility

$$\begin{aligned} \text{util}_v(2, s_{-v}) \quad = \quad & [\mathtt{c} \cdot |s|_1^v & + \quad \mathtt{b} \cdot |s|_2^v] \\ + \quad & [\mathtt{g} \cdot (|s|_1 - |s|_1^v) & + \quad \mathtt{f} \cdot (|s|_2 - 1 - |s|_2^v)] \ . \end{aligned}$$

Note that symmetric 2×2 games are known to be exact potential games [6, 40], and the potential is given as follows:

$$\Phi^c = \begin{pmatrix} \mathtt{a} - \mathtt{c} & 0 \\ 0 & \mathtt{b} - \mathtt{d} \end{pmatrix} \qquad \Phi^d = \begin{pmatrix} \mathtt{e} - \mathtt{g} & 0 \\ 0 & \mathtt{f} - \mathtt{h} \end{pmatrix}.$$

Here the potential for two players playing strategies i and j respectively, where $i, j \in \{1,2\}$, is $\Phi^c(i,j)$ for Γ^c and $\Phi^d(i,j)$ for Γ^d.

Proposition 1. *Every local interaction game has an exact potential function $\Phi(s)$ defined as sum of the corresponding potentials Φ^c and Φ^d of the symmetric 2×2 games.*

Proof. Consider an arbitrary state s and a player v changing his strategy.

$$\begin{aligned} & \text{util}(1, s_{-v}) - \text{util}(2, s_{-v}) \\ = \ & \sum_{\substack{(u,v)\in E \\ s_u=1}} \mathtt{a} + \sum_{\substack{(u,v)\in E \\ s_u=2}} \mathtt{d} + \sum_{\substack{(u,v)\notin E \\ s_u=1}} \mathtt{e} + \sum_{\substack{(u,v)\notin E \\ s_u=2}} \mathtt{h} - \left(\sum_{\substack{(u,v)\in E \\ s_u=1}} \mathtt{c} + \sum_{\substack{(u,v)\in E \\ s_u=2}} \mathtt{b} + \sum_{\substack{(u,v)\notin E \\ s_u=1}} \mathtt{g} + \sum_{\substack{(u,v)\notin E \\ s_u=2}} \mathtt{f} \right) \\ = \ & \sum_{\substack{(u,v)\in E \\ s_u=1}} ((\mathtt{a} - \mathtt{c}) - 0) + \sum_{\substack{(u,v)\in E \\ s_u=2}} (0 - (\mathtt{b} - \mathtt{d})) + \sum_{\substack{(u,v)\notin E \\ s_u=1}} ((\mathtt{e} - \mathtt{g}) - 0) + \sum_{\substack{(u,v)\notin E \\ s_u=2}} (0 - (\mathtt{f} - \mathtt{h})) \\ = \ & \sum_{(u,v)\in E} \Phi^c(1, s_u) - \Phi^c(2, s_u) + \sum_{(u,v)\notin E} \Phi^d(1, s_u) - \Phi^d(2, s_u) \\ = \ & \Phi(1, s_{-v}) - \Phi(2, s_{-v}) \ . \end{aligned}$$

□

2 Sequential Dynamics

In this section, our goal is to examine the duration of sequential iterative better-response dynamics. We provide an analysis of the potential function, which yield polynomial convergence times in 2-type interaction games.

Theorem 1. *For every 2-type interaction game every sequence of better-response moves from any initial state terminates in a pure Nash equilibrium after at most* $(n+1)(m+1)^2$ *steps.*

Proof. Our proof relies on a more insightful characterization for the potential function. Looking carefully at the structure of the potential function, we can derive games in the doubly symmetric form described by Figure 2, which have the same potential function. This implies that they also have the same payoff differences for the players and, in particular, the same Nash equilibria. We use $\mathtt{A} = \mathtt{a} - \mathtt{c}$, $\mathtt{B} = \mathtt{b} - \mathtt{d}$, $\mathtt{E} = \mathtt{e} - \mathtt{g}$, and $\mathtt{F} = \mathtt{f} - \mathtt{h}$. A trivial calculation (see, e.g., chapters 1 and 2 of [39]) shows that the new game exhibits the same potential functions as the original game. Note that this game is not equivalent in terms of social welfare, as we alter the total payoffs in some of the states.

Γ^c	1	2
1	A, A	0,0
2	0,0	B, B

Γ^d	1	2
1	E, E	0,0
2	0,0	F, F

Figure 2: Payoffs in games transformed to be doubly symmetric.

We analyze the potential function more closely and denote by $\mathtt{S} = \mathtt{A} + \mathtt{B}$, $\mathtt{T} = \mathtt{E} + \mathtt{F}$, $\Delta\mathtt{A} = \mathtt{A} - \mathtt{E}$, and $\Delta\mathtt{B} = \mathtt{B} - \mathtt{F}$. The potential function of Γ_p is

$$\begin{aligned}
\Phi(s) &= \sum_{v \in V_1(s)} \mathtt{A} \cdot |s|_1^v + \mathtt{E} \cdot (|s|_1 - 1 - |s|_1^v) + \sum_{v \in V_2(s)} \mathtt{B} \cdot |s|_2^v + \mathtt{F} \cdot (|s|_2 - 1 - |s|_2^v) \\
&= \mathtt{E} \cdot |s|_1 \cdot (|s|_1 - 1) + \mathtt{F} \cdot |s|_2 \cdot (|s|_2 - 1) + \sum_{v \in V_1(s)} \Delta\mathtt{A} \cdot |s|_1^v + \sum_{v \in V_2(s)} \Delta\mathtt{B} \cdot |s|_2^v \\
&= \mathtt{E} \cdot |s|_1 \cdot (|s|_1 - 1) + \mathtt{F} \cdot (n - |s|_1)(n - |s|_1 - 1) - \sum_{v \in V_1(s)} \Delta\mathtt{A} \cdot |s|_2^v + \sum_{v \in V_1(s)} \Delta\mathtt{A} \cdot \deg(v) \\
&\quad - \sum_{v \in V_2(s)} \Delta\mathtt{B} \cdot |s|_1^v + \sum_{v \in V_2(s)} \Delta\mathtt{B} \cdot \deg(v) \\
&= \mathtt{F}n(n-1) + \mathtt{T} \cdot (|s|_1)^2 - (2\mathtt{F}(n-1) + \mathtt{T}) \cdot |s|_1 + (\mathtt{T} - \mathtt{S}) \cdot m_{12}(s) \\
&\quad + \sum_{v \in V_1(s)} \Delta\mathtt{A} \cdot \deg(v) + \sum_{v \in V_2(s)} \Delta\mathtt{B} \cdot \deg(v) \\
&= \mathtt{F}n(n-1) + \mathtt{T} \cdot (|s|_1)^2 - (2\mathtt{F}(n-1) + \mathtt{T}) \cdot |s|_1 + (\mathtt{T} - \mathtt{S}) \cdot m_{12}(s) \\
&\quad + 2\Delta\mathtt{B}m + (\Delta\mathtt{A} - \Delta\mathtt{B}) \sum_{v \in V_1(s)} \deg(v) \; .
\end{aligned}$$

It is possible to drop the constant terms $\mathtt{F}n(n-1) + 2\Delta\mathtt{B}m$ from every state and derive an equivalent potential function

$$\begin{aligned}
\Psi(s) &= \Phi(s) - \mathtt{F}n(n-1) - 2\Delta\mathtt{B}m && (1) \\
&= \mathtt{T} \cdot (|s|_1)^2 - (2\mathtt{F}(n-1) + \mathtt{T}) \cdot |s|_1 + (\mathtt{T} - \mathtt{S}) \cdot m_{12}(s) + (\Delta\mathtt{A} - \Delta\mathtt{B}) \sum_{v \in V_1(s)} \deg(v). && (2)
\end{aligned}$$

For the proof of the theorem we observe that Ψ depends—in addition to the payoffs—only on three parameters: the number $|s|_1$ of players playing strategy 1, their degrees $\sum_{v \in V_1(s)} \deg(v)$ and the cut size $m_{12}(s)$. Observe that $|s|_1$ can range from 0 to n, which constitutes the factor $n+1$ in the guarantee. $m_{12}(s)$ and $\sum_{v \in V_1(s)} \deg(v)$ can take at most $m+1$ different values each. Hence, the total number of possible combinations for these parameters yields a total of $(n+1)(m+1)^2$ different values for Ψ and Φ. As each better-response iteration must strictly increase Φ in each step, every such sequence takes at most this number of iterations to reach a local optimum of Φ, from any starting state. This proves the theorem. $\square$

The main technique in the previous proof is transforming any game to an equivalent doubly symmetric game with only four different payoff values. The main outcome of this is the function Ψ in Equation (2), a potential function for our original game with an insightful representation.

The basis of the previous proof is a simple argument that can be applied somewhat more generally. Suppose every pair of players plays the same, exact potential game, each player can pick his strategy only once for all games, and the payoffs he receives are summed up. Then the whole game has an exact potential function. Consider a local interaction game in which each pair of players plays a $k \times k$ potential game with constant k. We can classify edges into $O(k^2)$ classes depending on the current state of the game on the edge. This yields only a polynomial number of different combinations and potential values. The same holds if we generalize 2-type interaction games to a constant number of different $k \times k$ potential games with constant k. On the other hand, if we allow a different game on each edge, then even with $k = 2$ this more general class of games includes the party affiliation or weighted MaxCut game as a special case, for which finding a pure Nash equilibrium is known to be PLS-hard, and in which there are games and starting states from which every sequence of better-response steps is of exponential length [16, 28]. Similarly, the class of local interaction games with $k \times k$ games and $k \leq n$ extends weighted MaxCut games, as we can simultaneously embed a MaxCut game into payoff matrix and graph structure for a subgraph of $k/2$ nodes. Thus, for $k = \Omega(n)$ strategies, the class again includes MaxCut games in which convergence time is necessarily exponential. Thus, in these natural extensions of our model, a similar result as in Theorem 1 is not possible. A deeper characterization along these lines is left for future work.

3 Concurrent Dynamics

In this section we consider round-based concurrent dynamics, in which in each round all players simultaneously update their strategy choices. A simple approach, which is considered frequently in the area of information diffusion in networks [33], is to allow all players to simultaneously play their best responses to the current state of the game. This approach converges rapidly if all players have dominant strategies. In fact, we would reach the dominant strategy equilibrium after the first round, which speeds up the convergence time by a factor of n over the sequential process considered previously. One might think that concurrent dynamics should always yield a significant speed-up due to the possibility of simultaneous updates. However, due to the absence of global coordination, these dynamics can easily get stuck in oscillations. The main design challenge proves to be to avoid oscillation and to obtain reasonable convergence times. In order to do this we follow the idea of [17] and design a policy in order to increase the potential function in expectation in each round. The challenge here is to enlarge migration probabilities to converge quickly, yet to guarantee potential increase in expectation.

To guarantee convergence we introduce the notion of inertia, i.e., we allow players with a certain probability to continue playing their strategy although there is currently a better alternative. For instance, suppose each player independently at random migrates to a better response with a probability less than 1. This allows for the construction of a Markov chain on the states, where migration probabilities of the players yield transition probabilities between states. Note that, due to inertia, with a possibly tiny probability the concurrent process can resemble any sequential better-response dynamics. Thus, the only absorbing states of the Markov chain are the pure Nash equilibria, to which the process must converge with probability 1 in the limit (see, e.g., [35]). The bounds on the convergence time that can be derived from this argument, however, are usually extremely large.

Subsequently, we analyze a protocol with migration probabilities proportional to the relative payoff increase. For technical reasons, we here assume that all payoffs are non-negative integer numbers, i.e., $\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}, \mathtt{e}, \mathtt{f}, \mathtt{g}, \mathtt{h} \in \mathbb{N}$. We first show that this protocol is unlikely to get stuck in oscillation. Afterwards, we consider several preprocessing steps to adjust the payoff values such that the incentives of players are preserved and convergence is obtained in expected polynomial time.

Algorithm 1 Relative Migration Protocol (RMP), repeatedly executed by all players in parallel.

For player v let $x \leftarrow s_v$ and $y \leftarrow 3 - x$.
if $\text{util}_v(y, s_{-v}) > \text{util}_v(x, s_{-v})$ **then**
  with probability

$$\mu_{xy} = \frac{1}{\lambda} \cdot \frac{\text{util}_v(y, s_{-v}) - \text{util}_v(x, s_{-v})}{\text{util}_v(y, s_{-v})}$$

  migrate from strategy x to y.
end if

In a state s a player v considers changing from strategy $x \in \{1, 2\}$ to strategy $y = 3 - x$ if $\text{util}_v(y, s_{-v}) - \text{util}_v(x, s_{-v}) > 0$. If this is the case, she migrates with migration probability that depends on her relative payoff increase (see the Relative Migration Protocol (RMP), Algorithm 1). Observe that to execute the RMP a player only needs to know its strategy and the strategy of its neighbors. If every player updates his strategy choices using the RMP, a new state s' evolves. We define a vector $\Delta s = (s'(v) - s(v))_{v \in V}$.

Lemma 1. *If $\mathtt{c} = \mathtt{d}$, $\mathtt{g} = \mathtt{h}$, and λ chosen sufficiently large, then as long as the 2-type interaction game is not in a Nash equilibrium, it holds that $\mathbb{E}[\Phi(s + \Delta s)] > \Phi(s)$.*

Say player v could improve his utility by switching to a new strategy. He decides to switch with a probability based on the action profile of his neighbors. At the same time as v changes strategy, his neighbors might do so as well. Thus this proof works by bounding the error in how much v expects to gain before switching versus how much v actually gains after switching.

Proof. For a state s and a vector Δs consider a player v. Let $x = s(v)$ denote v's current strategy and let $y = s(v) + \Delta s(v)$ denote v's strategy after migration. The change in v's utility after migration, assuming no other players change their strategy is denoted $\Delta\text{util}_v(s_{-v}) = \text{util}_v(y, s_{-v}) - \text{util}_v(x, s_{-v})$. Let the *virtual potential gain* be defined as

$$VPG(s, \Delta s) = \sum_{v \in V} \Delta\text{util}_v(s_{-v}).$$

The virtual potential gain simply sums all the presumed payoff increases of all players that chose to migrate. The real potential gain $\Phi(s + \Delta s) - \Phi(s)$ can be different if more than a single player moves. In this case the simultaneous migration of players u and v creates an error $F^{u,v}(s, \Delta s)$. Thus,

$$\Phi(s + \Delta s) - \Phi(s) = VPG(s, \Delta s) - \sum_{u,v \in V} F^{u,v}(s, \Delta s). \tag{3}$$

In order to show that $\mathbb{E}[\Phi(s+\Delta s)] - \Phi(s) > 0$, and conclude the proof of Lemma 1, we will relate expected virtual potential gain and expected error, which are the two terms on the right hand side of Equation (3). We denote by $\lambda^* = \gamma \cdot \frac{1}{2} \cdot \left(1 + \max\left\{\frac{\mathtt{a}}{\mathtt{b}}, \frac{\mathtt{b}}{\mathtt{a}}, \frac{\mathtt{e}}{\mathtt{f}}, \frac{\mathtt{f}}{\mathtt{e}}\right\}\right)$, where $\gamma > 1$ is a constant.

Lemma 2. *If $\mathtt{c} = \mathtt{d}$, $\mathtt{g} = \mathtt{h}$, then for $\lambda > \lambda^*$ it holds that $\mathbb{E}[\Phi(s + \Delta s)] - \Phi(s) \geq \frac{\gamma - 1}{\gamma} \cdot \mathbb{E}[VPG(s, \Delta s)]$.*

Proof. We will show that the error terms $\sum_{u,v \in V} \mathbb{E}[F^{u,v}(s, \Delta s)]$ are at most a constant fraction of $\mathbb{E}[VPG(s, \Delta s)]$, and the lemma will follow by taking the expectation of Eqn. (3). We will account the expected virtual potential gain partially to each pair of nodes $u, v \in V$, and thereby relate it to the expected error of the potential gain between u and v. For simplicity we drop the indices s, Δs and s_{-v}. Note that

$$\mathbb{E}[VPG] = \sum_{v \in V} \mu_{xy}^v \cdot \Delta \text{util}_v = \frac{1}{\lambda} \sum_{v \in V} \frac{(\Delta \text{util}_v)^2}{\text{util}_v(y)} = \frac{1}{\lambda} \sum_{v \in V} \Delta \text{util}_v \cdot R_v \ ,$$

where μ_{xy} is defined in Algorithm 1 and we use $R_v = \Delta\text{util}_v / \text{util}_v(y)$. We split this expected virtual potential gain into parts denoted $VPG^{u,v}$, which are accounted towards the pair (u,v) of players, for every $u \neq v$, $u,v \in V$. For a player v we account a fraction of his gain depending on the payoff that the game with player u contributes to $\text{util}_v(y)$.

The following analysis is done for a player v with $s(v) = x = 1$ that migrates to strategy $y = 2$ and pairs of neighboring players. The arguments can be repeated similarly for a switch from 2 to 1 and/or disconnected players. We first consider a neighbor u with $s(u) = 2$. For player v we account a fraction of

$$\frac{\mathtt{b}}{\text{util}_v(2)} \cdot \mu_{12}^v \cdot \Delta\text{util}_v = \mathtt{b} \cdot \frac{1}{\lambda} \cdot \left(\frac{\Delta\text{util}_v}{\text{util}_v(2)}\right)^2 = \frac{\mathtt{b}}{\lambda} \cdot R_v^2$$

of the expected virtual potential gain to the edge (u,v). Similarly, u has an incentive to change from strategy 2 to 1, and we account a fraction of

$$\frac{\mathtt{a}}{\text{util}_u(1)} \cdot \mu_{21}^u \cdot \Delta\text{util}_u = \mathtt{a} \cdot \frac{1}{\lambda} \cdot \left(\frac{\Delta\text{util}_u}{\text{util}_u(1)}\right)^2 = \frac{\mathtt{a}}{\lambda} \cdot R_u^2$$

of the expected virtual potential gain to (u,v). Thus, we have

$$\mathbb{E}[VPG^{u,v}] = \frac{1}{\lambda} \cdot (\mathtt{a}R_u^2 + \mathtt{b}R_v^2).$$

The expected error is calculated as follows. Player v presumes a change in payoff of $\mathtt{b} - \mathtt{c}$, player u presumes $\mathtt{a} - \mathtt{d} = \mathtt{a} - \mathtt{c}$. However, if both players migrate the combined change in potential is $(\mathtt{a} - \mathtt{c}) + (\mathtt{c} - \mathtt{a}) = 0$. Thus, the error is $\mathtt{a} + \mathtt{b} - 2\mathtt{c}$, and the expected error is

$$\mathbb{E}[F^{u,v}] = \mu_{21}^u \mu_{12}^v \cdot (\mathtt{a} + \mathtt{b} - 2\mathtt{c}) \leq \frac{\mathtt{a} + \mathtt{b}}{\lambda^2} \cdot R_u \cdot R_v \ .$$

Thus, the expected potential gain for the pair (u,v) is at least

$$\mathbb{E}[VP^{u,v}] - \mathbb{E}[F^{u,v}] \geq \frac{1}{\lambda}\left(\mathtt{b}R_v^2 + \mathtt{a}R_u^2 - \frac{(\mathtt{a}+\mathtt{b})}{\lambda} R_u R_v\right) \ .$$

This expression is strictly positive if we ensure that $(\mathtt{a}+\mathtt{b})/\lambda < 2\min\{\mathtt{a},\mathtt{b}\}$, which yields $\lambda > \frac{1}{2} \cdot \left(1 + \max\left\{\frac{\mathtt{a}}{\mathtt{b}}, \frac{\mathtt{b}}{\mathtt{a}}\right\}\right)$. By our choice of $\lambda > \lambda^*$ this is guaranteed and yields

$$\mathbb{E}[F^{u,v}] \leq \frac{1}{\gamma} \cdot \mathbb{E}[VPG^{u,v}] \ .$$

This proves the case for a neighbor u with $s(u) = 2$.

The case for a neighbor u with $s(u) = 1$ follows similarly. For player v we account a fraction of

$$\frac{\mathtt{c}}{\text{util}_v(2)} \cdot \mu_{12}^v \cdot \Delta\text{util}_v = \frac{\mathtt{c}}{\lambda} \cdot R_v^2$$

of his expected virtual potential gain to the edge (u,v). Similarly, u has an incentive to change from strategy 1 to 2, and we account a fraction of

$$\frac{\mathtt{c}}{\text{util}_u(2)} \cdot \mu_{12}^u \cdot \Delta\text{util}_u = \frac{\mathtt{c}}{\lambda} \cdot R_u^2$$

of the expected virtual potential gain to (u, v). Thus, we have

$$\mathbb{E}\left[VPG^{u,v}\right] = \frac{\mathtt{c}}{\lambda} \cdot (R_u^2 + R_v^2)$$

The expected error on edge (u, v) is calculated as follows. Each player presumes a change in payoff of $\mathtt{c} - \mathtt{a}$, however, if both players migrate, the change in potential is $\mathtt{c} - \mathtt{a} + \mathtt{b} - \mathtt{c} = \mathtt{b} - \mathtt{a}$. Thus, the error is $2(\mathtt{c} - \mathtt{a}) + (\mathtt{a} - \mathtt{b}) = 2\mathtt{c} - \mathtt{a} - \mathtt{b}$, and the expected error is

$$\mathbb{E}\left[F^{u,v}\right] = \mu_{12}^u \mu_{12}^v \cdot (2\mathtt{c} - \mathtt{a} - \mathtt{b}) \leq \frac{2\mathtt{c}}{\lambda^2} \cdot R_u \cdot R_v \ .$$

Due to the consistent factor $\mathtt{c}$ in this case we can actually argue with $\mathtt{b}, \mathtt{c} \geq 0$ and simple calculus that for any constant $\lambda > 1$ we have

$$\mathbb{E}\left[F^{u,v}\right] \leq \frac{1}{\lambda} \cdot \mathbb{E}\left[VPG^{u,v}\right] \leq \frac{1}{\gamma} \cdot \mathbb{E}\left[VPG^{u,v}\right] \ .$$

The same argument can be repeated for all pairs of players and all possible strategy constellations. Finally, we see that

$$\sum_{u,v \in V} \mathbb{E}\left[F^{u,v}\right] \leq \frac{1}{\gamma} \cdot \mathbb{E}\left[VPG\right]$$

This combined with Equation (3) proves Lemma 2. □

Note that, as long as at least one payoff value of $\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}$ is strictly positive, we make a strictly positive increase in the potential function whenever a player moves. This proves Lemma 1. □

In the following we will adjust local interaction games such that we preserve the incentives of players and the dynamics resulting from the RMP converge to a Nash equilibrium in expected polynomial time. We first turn the games into doubly symmetric games of the form in Figure 2. We then use the fact that for any such local interaction game we can find a game that preserves all player preferences and has $\mathtt{A}, \mathtt{B} \in [-2n^2, 2n^2]$, which we prove below. Finally, adjusting doubly symmetric games of Fig. 2 to ensure that all payoffs are positive is straightforward by adding $2\max\{|\mathtt{A}|, |\mathtt{B}|\}$ to every payoff value. Note that this adjustment at most triples the maximum absolute value of all payoffs. When we run the RMP in the equivalent game with payoffs adjusted in this way, we can guarantee polynomial convergence time. We refer to this as the *perturbed RMP*.

Note that perturbed RMP induces the same preference relation for the players as the original game. That is, a player v prefers strategy 1 over 2 for given strategies s_{-v} in the original game if and only if he does so in the perturbed game. However, we change the actual payoff values for the players to be of polynomial size. This only has an effect on the *amount* by which players prefer states, and results in altered migration probabilities. For instance, it always ensures that they are at least polynomially large (apart from being 0).

Theorem 2. *For local interaction games the dynamics resulting from the perturbed RMP converge to a Nash equilibrium in expected polynomial time.*

Proof. We first observe that we can always replace payoff values by numbers in $O(n^2)$ that yield the same player incentives. Then we show that this adjustment results in polynomial convergence time.

Let us consider the game Γ_p of the form in Fig. 2. Note that if $\mathtt{A} \geq 0 \geq \mathtt{B}$ or $\mathtt{B} \geq 0 \geq \mathtt{A}$, the game has a weakly dominant strategy and we get an equivalent game with $\mathtt{A}, \mathtt{B} \in \{1, 0, -1\}$. We here treat the case that $\mathtt{A}, \mathtt{B} > 0$, the completely negative case is similar.

Consider any player $v \in V$. The number of neighbors of v playing strategy 1 or 2 define the payoffs and the preferences of v. In particular, for every player v of degree $\deg(v)$ the utility function is given by $|s|_1^v\mathtt{A}$ or $(\deg(v) - |s|_1^v)\mathtt{B}$ if s_v is 1 or 2, respectively. For fixed values $\mathtt{A}, \mathtt{B}$ the preference for strategy 1 or 2 for

v depends only on the number of neighbors playing strategy 1. Consider the preference function $\text{pref}(x, y)$ that depends on $\mathtt{A}$ and $\mathtt{B}$ and is given by

$$\text{pref}(x,y) = \begin{cases} 0 & \text{if } x\mathtt{A} = y\mathtt{B} \\ 1 & \text{if } x\mathtt{A} < y\mathtt{B} \\ 2 & \text{if } x\mathtt{A} > y\mathtt{B} \end{cases} .$$

Note that for every player v and every state s of every local interaction game with payoffs $\mathtt{A}$ and $\mathtt{B}$, strategy $\text{pref}(|s|_1^v, \deg(v) - |s|_1^v)$ is the strategy preferred by player v in state s, where strategy 0 means that v prefers to stick to s_v.

We will now define integer values $\mathtt{A}', \mathtt{B}'$ with $1 \leq \mathtt{A}', \mathtt{B}' \leq 2n^2$ such that the preference function $\text{pref}'(x, y)$ resulting from $\mathtt{A}', \mathtt{B}'$ has $\text{pref}'(x, y) = \text{pref}(x, y)$ for all $x, y \in \{0, 1, \ldots, n\}$. This implies that for every local interaction game with $\mathtt{A}, \mathtt{B} > 0$, there is an equivalent local interaction game with $\mathtt{A}', \mathtt{B}' \in O(n^2)$ in the sense that for every player v and every state s, the incentives of v in s are the same in both games. Thus, in particular, the better-response dynamics and the set of pure Nash equilibria is the same in both games. Obviously, however, the exact utility values of players for the states may differ.

Note that we always have $\text{pref}(0, 0) = \text{pref}'(0, 0) = 0$. Furthermore, as $\mathtt{A}, \mathtt{B} > 0$ by assumption, we have $\text{pref}(0, y) = \text{pref}'(0, y) = 2$ and $\text{pref}(x, 0) = \text{pref}'(x, 0) = 1$ whenever $\mathtt{A}', \mathtt{B}' > 0$ and $x, y \in \{1, \ldots, n\}$. For the remaining values, we rearrange the definition of pref and use the positivity of x and y to observe

$$\text{pref}(x,y) = \left\{ \begin{array}{c} 0 \\ 1 \\ 2 \end{array} \right\} \Longleftrightarrow \left\{ \begin{array}{c} \mathtt{A}/\mathtt{B} = y/x \\ \mathtt{A}/\mathtt{B} > y/x \\ \mathtt{A}/\mathtt{B} < y/x \end{array} \right\} .$$

Hence, using

$$B^+ = \min_{x,y \in \{1,\ldots,n\}} \left\{ \frac{y}{x} \mid \text{pref}(x,y) = 2 \text{ or } \text{pref}(x,y) = 0 \right\}$$

and

$$B^- = \max_{x,y \in \{1,\ldots,n\}} \left\{ \frac{y}{x} \mid \text{pref}(x,y) = 1 \text{ or } \text{pref}(x,y) = 0 \right\}$$

implies $B^- \leq \mathtt{A}/\mathtt{B} \leq B^+$. By definition, one can express $B^- = b_1^-/b_2^-$ and $B^+ = b_1^+/b_2^+$ for values $b_1^+, b_1^-, b_2^+, b_2^- \in \{1, \ldots, n\}$. Now if $\mathtt{A}/\mathtt{B} = B^-$, then pick $\mathtt{A}' = b_1^-$ and $\mathtt{B}' = b_2^-$. Similarly, pick $\mathtt{A}' = b_1^+$ and $\mathtt{B}' = b_2^+$ if $\mathtt{A}/\mathtt{B} = B^+$. Otherwise, pick $\mathtt{A}', \mathtt{B}' \in \{1, \ldots, 2n^2\}$ such that

$$\begin{aligned} \mathtt{A}' &= b_1^- b_2^+ + b_1^+ b_2^- \\ \mathtt{B}' &= 2 b_2^- b_2^+ . \end{aligned}$$

Then $\mathtt{A}'/\mathtt{B}' \in (B^-, B^+)$. With this choice it is straightforward to observe that for all $x, y \in \{1, \ldots, n\}$

$$\text{pref}(x,y) = \left\{ \begin{array}{c} 0 \\ 1 \\ 2 \end{array} \right\} \Longleftrightarrow \left\{ \begin{array}{c} \mathtt{A}/\mathtt{B} = y/x \\ \mathtt{A}/\mathtt{B} > y/x \\ \mathtt{A}/\mathtt{B} < y/x \end{array} \right\} \Longleftrightarrow \left\{ \begin{array}{c} \mathtt{A}'/\mathtt{B}' = y/x \\ \mathtt{A}'/\mathtt{B}' > y/x \\ \mathtt{A}'/\mathtt{B}' < y/x \end{array} \right\} \Longleftrightarrow \text{pref}'(x,y) = \left\{ \begin{array}{c} 0 \\ 1 \\ 2 \end{array} \right\} .$$

We have seen that we can obtain an equivalent game with the same preferences for every player and payoffs with absolute values in $O(n^2)$. Hence, as observed above, we can further adjust this to an equivalent game with non-negative payoffs bounded by $O(n^2)$. The perturbed RMP is the RMP in this adjusted game. Here we can observe $\lambda^* \in O(n^2)$, i.e., the expected virtual potential gain in each round of the dynamics is in $\Omega(1/n^4)$ and so is the expected potential increase. Examining the potential for local interaction games with payoffs in $O(n^2)$ reveals that the maximum value of the potential is bounded by $O(n^4)$. The expected time to reach a state of maximum potential is thus bounded by $O(n^8)$. This shows that the perturbed RMP converges to a Nash equilibrium in expected polynomial time. □

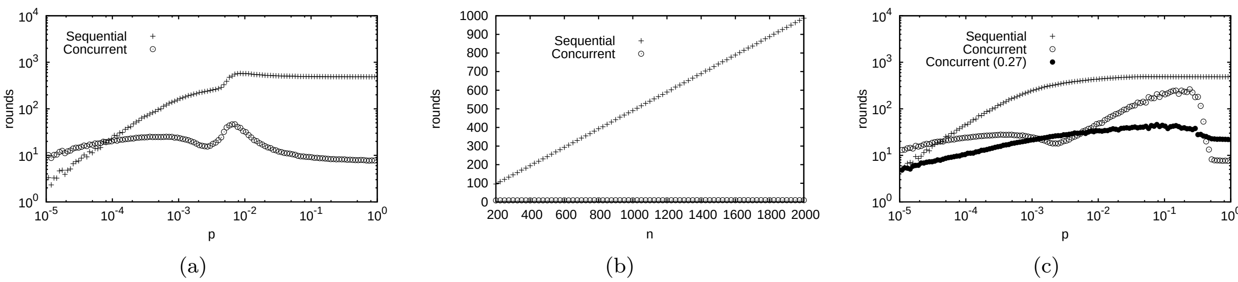

Figure 3: Running times of sequential and concurrent dynamics. (a) Coordination games on $G_{n,p}$ with $n = 1000$ and varying p. RMP dynamics are significantly faster for $p \geq \frac{1}{n^c}$ with $c < 1$. (b) Coordination games on $G_{n,p}$ with $p = \log^{-1}(n)$ and varying n. Sequential running times increase linearly, the times for RMP dynamics remain almost constant. (c) Coordination games on random unit disk graphs with $n = 1000$ and varying radius r. Running times of RMP dynamics are faster than sequential dynamics. Rapid convergence is achieved for concurrent dynamics with constant migration probability $\mu = 0.27$.

The previous theorem holds for local interaction games. For the more general class of 2-type interaction games, it is straightforward to argue that if an 2-type interaction game has payoffs polynomial in n, i.e., in $O(n^k)$ for some constant k, then $\lambda^* \in O(n^k)$ and the perturbed RMP yields an expected potential increase in each iteration that is in $\Omega(n^{-(2k+1)})$. In this case the maximum potential value is in $O(n^{k+2})$, which directly yields the following corollary.

Corollary 1. *For 2-type interaction games with payoffs bounded by $O(n^k)$ with a constant k the dynamics resulting from the perturbed RMP converge to a Nash equilibrium in expected polynomial time.*

For 2-type interaction games with arbitrary payoffs that are not necessarily bounded by $O(n^k)$, we conjecture that the results of Theorem 2 still hold. However, constructing a similar reduction as in Theorem 2 is quite a tedious approach. A formal proof of this conjecture is an interesting open problem.

4 Comparison of Convergence Times

The bound on convergence times presented in previous sections hold in general for any 2-type interaction game. However, there are significant differences between different types of games. We will exhibit these differences experimentally using the simpler local interaction games. In dominant strategy games concurrent dynamics have an obvious advantage, because there is no error when allowing players to migrate. In particular, by appropriately adjusting payoffs to 0 and 1 we can easily ensure that in the RMP every player migrates with arbitrarily large probability to the dominant strategy.

If there is no (weakly) dominant strategy, the game Γ_p is either a coordination game with $\mathtt{A}, \mathtt{B} > 0$, or an anti-coordination game with $\mathtt{A}, \mathtt{B} < 0$. For simplicity we restrict to *elementary* games, in which $\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d} \in \{0, 1\}$. For such games it is possible to show a time bound of $O(n^2)$ for sequential dynamics, and of $O(n^3)$ for concurrent dynamics resulting from the RMP, which will will call *RMP dynamics*.

4.1 Coordination Games

First we consider elementary coordination games with $\mathtt{a} = \mathtt{b} = 1$ and $\mathtt{c} = \mathtt{d} = 0$. The worst-case upper bound for the convergence time of RMP dynamics is a factor of $\Theta(n)$ larger than the bound for sequential dynamics. It is possible to design a game matching this difference, i.e., a game in which the RMP dynamics are a factor of $\Omega(n)$ slower than any sequential better-response dynamics. In this game, no concurrent better-response dynamics can be faster than *any* sequential dynamics.

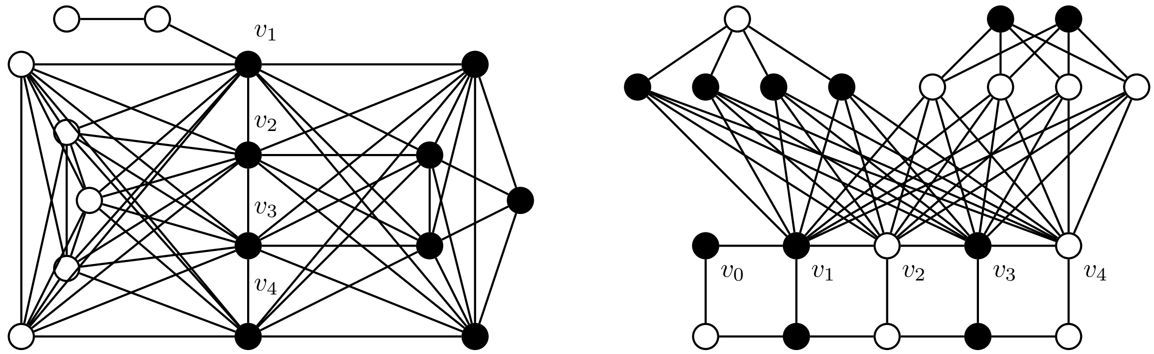

(a) (b)

Figure 4: Elementary (a) coordination and (b) anti-coordination games establishing a worst-case lower bound on the running time of concurrent dynamics with $k = 4$. Filled vertices play strategy 1, empty vertices play strategy 2. Starting with v players along the (a) middle and (b) upper line will switch to the opposite strategy one at a time.

Proposition 2. *For every $k \geq 1$ there is an elementary coordination game with $n = 3k+4$ players, in which every sequential better-response dynamics take exactly k steps and every concurrent dynamics take at least k steps. In particular, the RMP dynamics converge in $\Omega(kn)$ steps in expectation.*

Proof. Consider a game of the class depicted in Fig. 4(a). In this game we have a line of vertices $v_1, \ldots, v_k$, and all players currently play strategy 1. In addition, there are two cliques of size $k+1$, in one clique all players play 1, whereas in the other clique all players play 2. Finally, there are two additional vertices that play strategy 2. Each vertex on the line is connected to all vertices of the clique that plays 2 and to $k-1$ vertices of the clique that plays 1. Starting with v_1, in each round i the only player that wants to switch is player v_i from 1 to 2. Obviously, even if all players are given the possibility to jump, only this one player will migrate. For the RMP dynamics we observe that the relative improvement in payoff for the migrating player is in $\Theta(1/n)$ in each round, so it takes an expected number of $\Theta(n)$ rounds until one player migrates. □

This game reveals a fundamental dilemma for concurrent dynamics. On the one hand, it is possible to match the speed of sequential dynamics only if we let each player migrate with a large probability. On the other hand, frequent migration can yield long-lasting oscillations in other games. For the task of designing protocols with guaranteed rapid convergence, this problem is certainly critical. However, the bad properties are mainly due to the adversarial construction with an inherent partition into two parts that are intrinsically stable attached to opposite strategies.

We contrast these worst-case results with the average-case behavior on random graphs generated according to the $G_{n,p}$ model, as is done in the work of [29, 31]. We will observe similar behavior also on random unit disk graphs below. It turns out that in these games aggressive concurrent dynamics can make very rapid progress in initial stages.

Theorem 3. *Let $0 \leq c < 1$ be a constant and $\frac{1}{n^c} \leq p \leq \frac{1}{2}$, and let G be generated via $G_{n,p}$. Consider a state in the elementary coordination game with at least $(1/2+\delta)n$ nodes playing strategy 1 and at most $(1/2-\delta)$ nodes playing strategy 2, where $1/2 \geq \delta \geq 0$ is a constant. After 1 round of concurrent best-response dynamics all but $o(n)$ nodes will be playing strategy 1.*

Proof. For convenience let us define the S-degree of a vertex as follows.

Definition 2. *For a graph $G = (V, E)$ and $S \subseteq V$, the S-degree of a vertex $v \in V$ is* $\deg(v, S) = |\{u \in V \mid (u, v) \in E, u \in S\}|$.

The theorem follows by showing that a large number of vertices have similar neighborhoods. This idea is formalized in the following lemma and is similar to an argument given in [31].

Lemma 3. *Let $0 \leq c < 1$ be a constant and $\frac{1}{n^c} \leq p \leq \frac{1}{2}$. Let $S \subseteq V$ be such that $|S| \geq \tau n$ for some constant $\tau > 0$. For any constant $1/6 \geq \epsilon > 0$ the number of nodes that have S-degree outside the range $(1 \pm \epsilon)p\tau n$ is at most $o(n)$.*

Proof. The main vehicle for proving this lemma is Theorem 2.14 of Bollobas [8]. First we check that the necessary assumptions of that theorem are met: $\frac{6\log(n)}{\epsilon^2 p} \leq \frac{6\log(n)n^c}{\epsilon^2} \leq \tau n \leq |S|$. The inequalities hold because $\frac{1}{n^c} \leq p$ and $c < 1$. Thus for the set of nodes with a significantly unexpected number of neighbors $Z_S = \{z \in V \setminus S : ||\Gamma(z) \cap S| - p\tau n| \geq \epsilon p \tau n\}$ the theorem gives us with $\frac{1}{n^c} \leq p$ and $c < 1$ that

$$|Z_S| \leq \frac{12\log(n)}{\epsilon^2 p} \leq \frac{12\log(n)n^c}{\epsilon^2} = o(n)$$

□

We now apply the lemma to show our theorem. Since there are at least $(1/2 + \delta)n$ nodes playing 1, by Lemma 3 the number of nodes playing 2 with V_1-degree outside the range $(1/2 + \delta)(1 \pm \epsilon)pn$ is at most $o(n)$. Similarly, since there are at most $(1/2 - \delta)n$ nodes playing 2, by Lemma 3 the number of nodes with V_2-degree outside the range $(1/2 - \delta)(1 \pm \epsilon)pn$ is at most $o(n)$. We can choose ϵ small enough such that $(1/2 + \delta)(1 - \epsilon)pn > (1/2 - \delta)(1 + \epsilon)pn$. Thus, all but $o(n)$ of the nodes playing strategy 2 have more neighbors playing strategy 1 than strategy 2, and the best response for all but $o(n)$ of the nodes playing 2 is to switch to 1. By an analogous argument we can show that all but $o(n)$ of the nodes playing strategy 1 have more neighbors that play strategy 1 than 2. Thus all but $o(n)$ of the nodes playing 1 will have the best response of continuing to play 1. □

If the dynamics are sequential instead of concurrent, one can show by a similar argument to the above that after n rounds all but $o(n)$ nodes will be playing strategy 1.

Next, we show a number of experimental results in Fig. 3. For each value of n and p we generated 10 random graphs, and on each random graph we chose 25 starting states uniformly at random. From each starting state we initiated 25 runs of RMP dynamics. For the sequential dynamics we deterministically chose in each round one player that yields the largest payoff gain. The constant λ was set to $\lambda = 1.1$ throughout. Fig. 3(a) shows the average number of rounds for $n = 1000$ and p increasing exponentially between 10^{-5} and 1. When the large component forms (around $p = 0.005$) the sequential times are close to $n/2$, while the RMP dynamics converge rapidly in a constant number of runs.

Although Theorem 3 does not directly bound the convergence time to Nash equilibria, it provides the main intuition for the explanation of the results. After random initialization there are close to $n/2$ players playing each strategy. Afterwards, due to similar neighborhoods and coordination structure of the game, nearly all players accumulate on one strategy. Although this does not happen in one step, it still occurs quite rapidly, as each player migrating to a predominant strategy increases the probability for others to follow. Thus, in essence the behavior of the RMP dynamics is characterized by the insights from Theorem 3.

The intuition follows similarly for the sequential case, see Fig. 3(b). It depicts running times on graphs with increasing n and $p = \log^{-1}(n)$. Observe that RMP dynamics yield rapid convergence times that increase only very slightly. Sequential dynamics need roughly $\Theta(n)$ rounds until a Nash equilibrium is reached.

4.2 Anti-Coordination Games

The elementary anti-coordination game is the MaxCut game with $\mathtt{a} = \mathtt{b} = 0$ and $\mathtt{c} = \mathtt{d} = 1$. For this game the worst-case results are similar to the coordination case. More specifically, RMP dynamics can be a factor of $\Omega(n)$ slower than any sequential better-response dynamics, and the game reveals that every concurrent dynamics are at least as slow as any sequential dynamics.

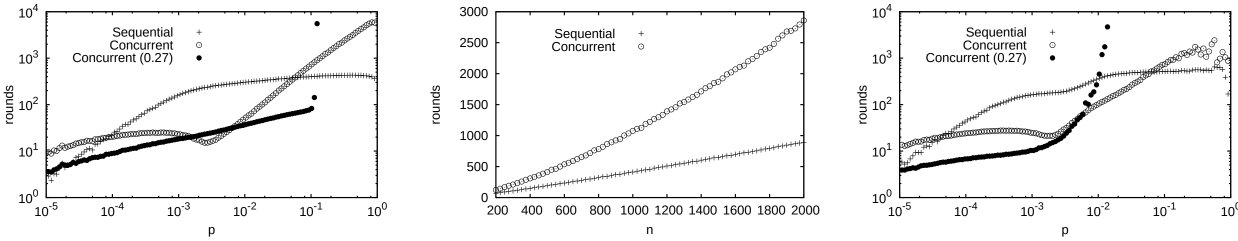

(a) (b) (c)

Figure 5: (a) Anti-coordination games on $G_{n,p}$ with $n = 1000$ and varying p. Running times of RMP dynamics increase linearly for $p \geq \frac{1}{n^c}$ with $c < 1$. For concurrent migration with probability $\mu = 0.27$, oscillation becomes dominant abruptly. (b) Anti-coordination games on $G_{n,p}$ with $p = \log^{-1}(n)$ and varying n. Sequential running times increase linearly, the times for RMP dynamics increase with roughly with $n \log n$. (c) Anti-coordination games on random unit disk graphs with $n = 1000$ and varying radius r. Running times of RMP dynamics increase with increasing node degree. For concurrent dynamics with constant migration probability, oscillation becomes dominant on dense graphs.

Proposition 3. *For every $k \geq 1$ there is an elementary anti-coordination game with $n = 4k + 4$ players, in which every sequential better-response dynamics take exactly k steps and every concurrent dynamics take at least k steps. In particular, the RMP dynamics converge in $\Omega(kn)$ steps in expectation.*

Proof. The class of games we consider is depicted in Fig. 4(b). There are two *pure sets* of k vertices each, one set playing strategy 1, one set playing strategy 2. Each set of vertices is connected to one stabilizing vertex of the opposite strategy. In addition, there are two lines of length $k + 1$, which are connected into a grid. Players on the second line play strategies alternatingly, starting with strategy 2. On the first line $v_0, \ldots, v_k$ the assignment is the same except for the leftmost vertex v_0, which is assigned to strategy 1. In addition, all vertices v_i with $i \geq 1$ are connected to all vertices from both pure sets. Starting in this configuration, player v_1 is the only vertex that wants to switch. In the following, in round i only player v_i will switch to the opposite strategy. Again, even if all players are given the possibility to jump, only one player will migrate. For the RMP dynamics we observe that the relative improvement in payoff for the migrating player is in $\Theta(1/n)$ in each round, so it takes an expected number of $\Theta(n)$ rounds until one player migrates. □

We complement this lower bound with experimental results in Fig. 5. Fig. 5(a) and 5(b) are generated using the same parameters as for Fig. 3(a) and 3(b), respectively. While for small p the behavior of both dynamics is similar to the coordination case, it changes when $p \geq \frac{1}{n^c}$ for $c < 1$ which corresponds to roughly $p \geq 10^{-2}$ in Fig. 5(a). Observe the linear increase in running time with growing p for the RMP dynamics, which for large p leads even to worse convergence times than for sequential dynamics. A linear dependence on p is also supported by Fig. 5(b), as here $p = \log^{-1}(n)$, and the time growth for the RMP dynamics is proportional to $n \log n$. In fact, the linear dependence is a result from the RMP dynamics being too passive. Unlike in the coordination case, players do not accumulate on one strategy choice. In most iterations there is no significant majority playing one strategy. Payoff differences remain small, so with degrees growing linear in p, migration probabilities μ^v drop to a level proportional to $1/p$. The expected time until a player migrates then grows linearly in p. This effect is present until p is very close to 1, in which case the convergence times of sequential dynamics drop to 0, as uniformly random initialization yields an almost stable profile. Furthermore, for almost complete graphs, the RMP dynamics yield a sequential process with high probability. This is because in very dense graphs almost all players have the same neighborhood and experience the same changes in payoff. The migration probabilities in the RMP dynamics of roughly $1/n$ are balanced by the $\Theta(n)$ players that are willing to migrate in each round, so there is a roughly constant number of player migrating in each round.

Large running times are due to the payoff-relative update rule of the RMP. With different choices it is possible to achieve much more rapid convergence. Fig. 5(a) also depicts the convergence times of concurrent dynamics on graphs with $n = 1000$ and varying p where all migration probabilities μ^v are chosen as a fixed value $\mu = 0.27$, other values yield similar results. The increased migration significantly decreases the expected running times below the sequential times. At some point, however, the dynamics rather abruptly hit an "oscillation barrier" and convergence times start growing exponentially. Characterizing this barrier and providing further analytical insights on suitable choices of migration probabilities in concurrent dynamics remains a fascinating open problem.

Finally, we note that the key observations hold similarly for the case of random unit-disk graphs, which are a popular model for interference in distributed networks. We generated graphs by placing n points uniformly at random in the unit square. An edge was created between two points if the distance under the maximum norm was at most r. For each graph we chose 25 starting states uniformly at random, and from each state we initiated 25 runs of the dynamics. We provide average running times in Figure 3(c) and 5(c).

5 Conclusion

We have studied distributed decision making in a fundamental class of network interaction games with various applications in distributed systems and social networking. Our results concern the convergence time of sequential and concurrent better-response dynamics. The analysis reveals polynomial convergence times for sequential dynamics in both local interaction games and 2-type interaction games. For concurrent dynamics resulting from the RMP there is polynomial convergence time in local interaction games, and in 2-type interaction games with polynomially bounded payoffs. In these games a local potential maximizer – i.e. a pure Nash equilibrium – can be obtained efficiently using distributed protocols, and thus efficient distributed decision making is possible. This stands in contrast to noisy better-response dynamics and global potential maximizers, which can even be NP-hard to compute (for instance, in MaxCut games). Even for coordination games, in which computation is trivial, noisy better-response dynamics can take exponential time to converge [37].

While our results establish a general upper bound, the actual convergence times differ significantly based on the type of interaction and the underlying network. Using experiments we have shed light on the influence of incentives and the degree of connectedness. More work is needed to obtain analytical characterizations for specific games and graph classes of interest.

Our dynamics result from the RMP using payoff-based probabilistic migration. While we use special assumptions (such as, e.g., non-negativity of payoffs) for defining the protocol, we can transform all local interaction games and 2-type interaction games (with polynomially bounded payoffs) into an equivalent representation that fits the needs of our protocol. It is an interesting open problem if we can also achieve fast convergence using a more direct approach, e.g., for migration probabilities chosen as a suitable constant, according to a logit response rule, or more general update rules [7]. It seems that the analysis of such rules would require to significantly extend the techniques we used in this paper.

Acknowledgment

The first author would like to thank Ulrik Brandes and Simon Fischer for discussion about the results in this work.

References

[1] Heiner Ackermann, Petra Berenbrink, Simon Fischer, and Martin Hoefer. Concurrent imitation dynamics in congestion games. In *Proc. 28th Symp. Principles of Distributed Computing (PODC)*, pages 63–72, 2009.

[2] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *J. ACM*, 55(6), 2008.

[3] Robert Aumann. Subjectivity and correlation in randomized strategies. *J. Mathematical Economics*, 1:67–96, 1974.

[4] Siegfried Berninghaus and Bodo Vogt. Network formation in symmetric 2×2 games. *Homo Oeconomicus*, 23(3/4):421–466, 2006.

[5] Avrim Blum and Yishay Mansour. Learning, regret minimization, and equilibria. In Nisan et al. [38], chapter 4.

[6] Lawrence Blume. The statistical mechanics of strategic interaction. *Games Econom. Behav.*, 5:387–424, 1993.

[7] Lawrence Blume. How noise matters. *Games Econom. Behav.*, 44(2):251–271, 2003.

[8] Béla Bollobás. *Random Graphs*. Cambridge University Press, second edition, 2001.

[9] Yann Bramoullé. Anti-coordination and social interactions. *Games Econom. Behav.*, 58(1):30–49, 2007.

[10] Yann Bramoullé, Dunia López-Pintado, Sanjeev Goyal, and Fernando Vega-Redondo. Network formation and anti-coordination games. *Int. J. Game Theory*, 33(1):1–19, 2004.

[11] William Brock and Steven Durlauf. Discrete choice with social interactions. *Rev. Econ. Stud.*, 68(2):235–260, 2001.

[12] George Christodoulou, Vahab Mirrokni, and Anastasios Sidiropoulos. Convergence and approximation in potential games. In *Proc. 23rd Symp. Theoretical Aspects of Computer Science (STACS)*, pages 349–360, 2006.

[13] Constantinos Daskalakis, Paul Goldberg, and Christos Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.

[14] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, 1993.

[15] Alex Fabrikant and Christos Papadimitriou:. The complexity of game dynamics: BGP oscillations, sink equilibria, and beyond. In *Proc. 19th Symp. Discrete Algorithms (SODA)*, pages 844–853, 2008.

[16] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proc. 36th Symp. Theory of Computing (STOC)*, pages 604–612, 2004.

[17] Simon Fischer. *Dynamic Selfish Routing*. PhD thesis, Lehrstuhl für Algorithmen und Komplexität, RWTH Aachen, 2007.

[18] Simon Fischer, Harald Räcke, and Berthold Vöcking. Fast convergence to Wardrop equilibria by adaptive sampling mehtods. *SIAM J. Comput.*, 39(8):3700–3735, 2010.

[19] David Foster and Rakesh Vohra. Calibrated learning and correlated equilibrium. *Games Econom. Behav.*, 21:40–55, 1997.

[20] Drew Fudenberg and David Levine. Consistency and cautious ficticious play. *J. Econ. Dynamics and Control*, 19:1065–1090, 1995.

[21] Drew Fudenberg and David Levine. *The theory of learning in games*. MIT Press, 1998.

[22] Ioannis Giotis and Venkatesan Guruswami. Correlation clustering with a fixed number of clusters. *Theory of Computing*, 2(1):249–266, 2006.

[23] Michel Goemans, Vahab Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *Proc. 46th Symp. Foundations of Computer Science (FOCS)*, pages 142–154, 2005.

[24] Sanjeev Goyal and Fernando Vega-Redondo. Network formation and social coordination. *Games Econom. Behav.*, 50:178–207, 2005.

[25] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

[26] Martin Hoefer and Siddharth Suri. Dynamics in network interaction games. In *Proc. 23rd Intl. Symp. Distributed Computing (DISC)*, pages 294–308, 2009.

[27] Matthew Jackson and Alison Watts. On the formation of interaction networks in social coordination games. *Games Econom. Behav.*, 41:265–291, 2002.

[28] David Johnson, Christos Papadimitriou, and Mihalis Yannakakis. How easy is local search? *J. Comput. Syst. Sci.*, 37:79–100, 1988.

[29] Sham Kakade, Michael Kearns, Luis Ortiz, Robin Pemantle, and Siddharth Suri. Economic properties of social networks. In *Advances in Neural Information Processing Systems (NIPS'04)*, 2004.

[30] Michael Kearns. Graphical games. In Nisan et al. [38], chapter 7.

[31] Michael Kearns and Siddharth Suri. Networks preserving evolutionary equilibria and the power of randomization. In *Proc. 7th Conf. Electronic Commerce (EC)*, pages 200–207, 2006.

[32] Michael Kearns and Jinsong Tan. Biased voting and the democratic primary problem. In *Proc. 4th Intl. Workshop Internet & Network Economics (WINE)*, pages 639–652, 2008.

[33] Jon Kleinberg. Cascading behavior in netwoks: Algorithmic and economic issues. In Nisan et al. [38], chapter 24.

[34] Jason Marden. *Learning in Large-Scale Games and Cooperative Control*. PhD thesis, UCLA, Los Angeles, 2007.

[35] Jason Marden, Gürdal Arslan, and Jeff Shamma. Joint strategy fictitious play with inertia for potential games. *IEEE Trans. Automatic Control*, 54(2):208–220, 2009.

[36] Dov Monderer and Lloyd Shapley. Potential games. *Games Econom. Behav.*, 14:1124–1143, 1996.

[37] Andrea Montanari and Amin Saberi. Convergence to equilibrium in local interaction games. In *Proc. 50th Symp. Foundations of Computer Science (FOCS)*, pages 303–312, 2009.

[38] Noam Nisan, Éva Tardos, Tim Roughgarden, and Vijay Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[39] Jörgen Weibull. *Evolutionary Game Theory*. MIT Press, 1995.

[40] H. Peyton Young. *Individual Strategy and social structure*. Princeton University Press, 1998.

[41] H. Peyton Young. *Strategic Learning and its Limits*. Oxford University Press, 2004.

[42] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proc. 20th Intl. Conf. Machine Learning (ICML'03)*, pages 928–936, 2003.