Local Dynamics in Bargaining Networks via Random-Turn Games

L. Elisa Celis[1]*, Nikhil R. Devanur[2], and Yuval Peres[2]

[1] University of Washington eceliscs.washington.edu
[2] Microsoft Research {nikdev, peres}@microsoft.com

Abstract. We present a new technique for analyzing the rate of convergence of local dynamics in bargaining networks. The technique reduces balancing in a bargaining network to optimal play in a random-turn game. We analyze this game using techniques from martingale and Markov chain theory. We obtain a tight polynomial bound on the rate of convergence for a nontrivial class of unweighted graphs (the previous known bound was exponential). Additionally, we show this technique extends naturally to many other graphs and dynamics.

1 Introduction

In a network bargaining game, nodes in a graph are involved in pairwise transactions with their neighbors. This type of game was introduced by Cook and Yamagishi [14] to capture the "power" of a node derived from its position in a network, and has also been used in economics to model two-sided markets [33, 31]. Recently these games have been analyzed from a computational point of view, first in a centralized model [23] and later in a distributed model [3]. Analyzing simple, local dynamics that converge quickly to an equilibrium in such games was an important open problem that attracted much interest [18, 22, 21].

We draw a connection between network bargaining games and random-turn games. Random-turn games are a well-studied class of two-player combinatorial games in which the outcome of a coin flip determines which player moves next [25, 24]. Combinatorial games can be represented as a game on a directed graph where players move a token along edges until one reaches their goal state. We transform the network bargaining game into an equivalent random-turn game which we can analyze using martingale techniques to obtain bounds on the rate of convergence. In particular, the convergence rate for the dynamics is related to the *absorption time* of the corresponding random-turn game.

We obtain a tight polynomial bound on the convergence rate for a variety of natural dynamics on a certain class of graphs. This class includes unweighted bipartite graphs with unique balanced outcomes, and the exposition is conducted in this setting for clarity. The previous bound known for any class of graphs (other than paths) was exponential.

* Research supported by a NASA Graduate Fellowship

Network Bargaining Game

A network bargaining game is defined on a weighted graph $G = (V', E')$ with $w' : E' \to \mathbb{R}_+$.[3] Every node in the graph is a player, and the weight of an edge represents the dollar amount available to be shared between the two adjacent players. However, each player is constrained to make at most one such sharing agreement. An *outcome* of this game is a *matching* in the graph $M \subseteq E$ and an *allocation* describing each player's profit, $f' : V' \to \mathbb{R}_+$ where for all $(uv) \in M$, we have $f'(u) + f'(v) = w'(uv)$, and for all unmatched $u \in V'$, $f'(u) = 0$.

We consider two notions of equilibrium in this game. The first (weaker) notion is that of a *stable* outcome: an outcome is stable if for all unmatched edges $(uv) \notin M$ we have $f'(u) + f'(v) \geq w'(uv)$, that is, no two adjacent players have incentive to deviate from their current matches. The second notion is that of a *balanced* outcome: an outcome is balanced if matched players divide the *surplus* equally amongst themselves. To be precise, let the *best alternate* of a node u be

$$\alpha_{f'}(u) := \max\{0, \max_{v:(uv)\in E'\setminus M} \{w'(uv) - f'(v)\}\},$$

i.e. the maximum profit a player could get from a neighbor she is not currently matched to. For every matched edge (uv) define the surplus as

$$s_{f'}(uv) = w'(uv) - (\alpha_{f'}(u) + \alpha_{f'}(v)).$$

An outcome is balanced if it is stable and for all matched edges (uv), $f'(u) = \alpha_{f'}(u) + s_{f'}(uv)/2$ and $f'(v) = \alpha_{f'}(v) + s_{f'}(uv)/2$, or equivalently, $f'(u) - \alpha_{f'}(u) = f'(v) - \alpha_{f'}(v)$. This can be seen as a generalization of Nash's bargaining solution for two players [28]. It is known that the following are equivalent: *(1)* a balanced outcome exists, *(2)* a stable outcome exists and *(3)* the matching polytope has no integrality gap [23].

Edge Balancing Dynamics

Balanced outcomes can be computed by centralized polynomial time algorithms [23], but the game is by nature distributed; individual players working on individual deals. An important open problem was to show there exist simple and natural *local dynamics* that converge quickly to a balanced outcome. We now define such dynamics with respect to a matching M and initial allocation f'.

For our dynamics, the matching M is fixed throughout. This may seem counter to the solution concept of a balanced outcome since the premise is the threat of switching partners. However, once such a threat is acknowledged, the players do not need to switch in order to bargain for their fair share. Moreover, there are distributed dynamics that find matchings [6, 32] which also have a bargaining flavor in their dynamics. One can imagine a two phase approach, where

[3] We reserve the notation (V, E) and w for a graph which will be used more prevalently in the random-turn game framework.

in the first phase the players find a matching and in the second find a balanced outcome with the matching fixed.

The allocations are updated synchronously, and the updates proceed in rounds. The allocation in round t is denoted by $B'_{f'}(u,t)$, the best alternatives by $\alpha_{f'}(u,t)$ and the surpluses by $s_{f'}(uv,t)$.[4] The initial allocation is $B'_{f'}(v,0) = f'(v)$. Synchronous Edge Balancing is defined by the following update rule: for all $u \in V'$, $(uv) \in M$ and $t \geq 1$,

$$B'_{f'}(u,t+1) \leftarrow \alpha_{f'}(u,t) + s_{f'}(uv,t)/2.$$

Thus, the allocation for the next round is determined by "balancing" each matched edge using the allocation in the current round.

We say that an allocation f' is *ε-close to balanced* if there exists a balanced outcome B' such that $|B'(v) - f'(v)| \leq \varepsilon$ for all v, i.e. we get ε-close to a balanced outcome. Note that this is stronger than a common alternate notion of ε-balanced where $|B'_{f'}(u,t+1) - B'_{f'}(u,t)| \leq \varepsilon$, i.e. each edge is locally balanced. We wish to show Synchronous Edge Balancing converges rapidly to a balanced outcome. This means that for all f', after polynomially many [5] time steps t, the allocation $B_{f'}(u,t)$ is ε-close to a balanced outcome.

Random-Turn Games

Every two-player game from Tic-Tac-Toe to Chess can be formalized as a combinatorial game on a directed graph where each turn consists of moving a token from one vertex to another along an edge [7]. Random-turn games are combinatorial games where the turns are determined by a coin flip.

We consider the following version in the main body of this paper: A Random-Turn Game consists of a directed graph $D = (V,E)$, payoff function $f : V \to [0,1]$, initial vertex v_0, and horizon $T \in \mathbb{N}$. The set V of game states contains two terminal states s and r and all payoff functions set $f(s) = 0$ and $f(r) = 1$. The game is played by Max and $Mini$ where Max's goal is to maximize the value of the end state, and $Mini$'s goal is to minimize it. Game play for horizon T is as follows: a token is initially placed at v_0 and at every step a fair coin is tossed to determine who gets to move the token. Max must always move to a predecessor of v and $Mini$ to a successor (as determined by the edge set E). We repeat until either T moves have been made, or we reach an absorbing state $\{s,r\}$. At the end of the game, $Mini$ pays Max \$$f(v)$ if the game terminates at node v. Since this is a full-information game, for any finite horizon, one can compute the optimal strategies for the two players. This defines a *value* of the game, which is the expected payoff for Max under optimal play.

Related Work and Motivation

Network bargaining games have a long history in two communities: sociology and game theory. In sociology, they are studied under the name *network exchange*

[4] The subscript f' may be dropped when it is clear from context.

[5] Where the polynomial is in $|V|, |E|$ and $1/\log(\varepsilon)$.

theory, where the goal is to understand the power of a node as a function of its position in the network (see the overview by Willer [35]). Network bargaining games as we define here were introduced by Cook and Yamagishi [14], who also introduced the notion of balanced outcomes. In fact, they also introduced local dynamics similar to what we consider in this paper, but without a theoretical analysis of the convergence of their dynamics. There have also been experimental results [13, 8] which validate the relevance and applicability of this work.

In game theory, the study of bargaining can be traced back to Nash's bargaining solution [28]. Many results in this field focus on two-sided markets, which naturally give rise to the bipartite version of the network bargaining game as was introduced by Shapley and Shubik [33]. This version, known as the *assignment game*, can also be viewed as the classic Gale-Shapley stable marriage problem [19] with the addition transferable utilities. Rochford [31] defined balanced outcomes for assignment games under the name symmetrically pairwise-bargained allocations. She also showed that they are the intersection of the core and the kernel, two common solution concepts in co-operative game theory. Other solution concepts such as the nucleolus [27] have also been considered. In fact, the computability of these solution concepts has been much studied [34, 29]. Other related models consider price setting as a result of a bargaining process [15].

Network bargaining games were introduced to the theoretical computer science community by Kleinberg and Tardos [23]. They gave a polynomial time algorithm to compute the set of balanced outcomes. Since then, there has been a flurry of activity: Azar, et al. [3] considered an asynchronous version of edge balancing dynamics and showed (exponential time) convergence. Other aspects of network bargaining have also been studied in the recent past [5, 10, 9, 4, 21].

We give the first polynomial time bound on local dynamics converging to a balanced outcome for any non-trivial class of graphs. The only polynomial time bound known previously was for paths. Moreover, the bounds are tight for a variety of dynamics. Independently and concurrently with our work, Kanoria, et al. [22] considered the same problem and showed convergence of a (different) dynamics to a balanced outcome. The dynamics they consider has the advantage that it does not need a matching to be known and fixed; rather, the dynamics also finds a matching. One drawback is that the outcome their process converges to is weaker (it is ε-balanced as opposed to ε-close to a balanced outcome). Additionally, the rate of convergence of their dynamics is weaker and of the form n^4/g^2 where no bounds on g are given. In fact, on many graphs where our result is tight, g could be zero.[6] Also independently and concurrently, Draief and Vojnovic [17] showed quadratic convergence of the edge balancing dynamics for the following graphs: a path, a cycle, a blossom and a bicycle. Faigle, Kern and Kuipers [18] also considered similar local dynamics for a more general class of games, but do not show bounds on the rate of convergence.

In general, analysis of the convergence of local dynamics to an equilibrium of a game is a common theme. Examples include analysis of random best response dynamics for the Gale-Shapley stable matching game [2, 19]. In fact, a

[6] For instance, this occurs on any unweighted even length path.

major philosophical hypothesis of algorithmic game theory [20, 12, 16] is that the existence of such dynamics is crucial to validate a solution concept.

Random-turn games are natural, and many variants have been analyzed [24, 25]. Most interestingly, a variant called the tug-of-war game has been found to be related to partial differential equations such as the infinity Laplacian and the p-Laplacian [30], due to which these games have received considerable attention [11, 37, 1].

Organization

In Section 2 we introduce our theorems, techniques and extensions. Section 3 contains a detailed analysis for unweighted bipartite graphs with unique balanced outcomes. We conclude and suggest future work in Section 4.

2 From a Bargaining Game to a Random-Turn Game

We now give a reduction from a network bargaining game to a random-turn game, the concept that lies at the heart of our results. We first restrict ourselves to unweighted bipartite graphs for clarity.

Consider a graph $G = (V', E')$ where $w'(uv) = 1$ for all $(uv) \in E'$ and V' is bipartitioned as $\{L, R\}$. Create a directed graph D as follows: let $D = (V, E)$ where V is the subset of matched vertices in L along with two special vertices, s and r. Let the set of vertices other than s and r be denoted by $\dot{V}$. Add an edge $(uv) \in E$ if $(M(u)v) \in E'$. Additionally, place an edge from s to all vertices in $\dot{V}$ and an edge from all vertices in $\dot{V}$ to r. Finally, add an (rv) edge in E if there exists an edge $(vu) \in E'$ where $u \notin M$. Similarly, add a (vs) edge if there is a $(M(v)u)$ edge with $u \notin M$. We also give an allocation $f : V \to [0, 1]$ on D, given the allocation f' on G. Define $f(v) = f'(v)$ if $v \in \dot{V}$, $f(s) = 0$ and $f(r) = 1$. See Figure 1 for an example of this reduction. Note that an allocation f' on G takes values between 0 and 1 since the edge weights all have weight 1. Thus, the definition of an allocation allows us to reconstruct f' from f, since $f'(M(v)) = 1 - f'(v)$ and $f(u) = 0$ when $u \notin M$.
The concepts (from the bargaining game described earlier) translate as follows.

- An allocation is *stable* if for all edges $(uv) \in E$, $f(u) \leq f(v)$.
- Let the *best predecessor* and *successor* of a node v be $v_f^+ = \arg\max_{u:(uv)\in E}\{f(u)\}$ and $v_f^- = \arg\min_{u:(vu)\in E}\{f(u)\}$ respectively. An allocation is *balanced* if it is stable, and for all vertices $v \in \dot{V}$, $f(v) = \frac{1}{2}(f(v_f^+) + f(v_f^-))$.
- Let the allocation in round t of Synchronous Edge Balancing be $B_f(v, t)$ where $B_f(v, 0) = f(v)$. Then, balancing is equivalent to $B_f(v, t + 1) = \frac{1}{2}(B_f(v^+, t) + B_f(v^-, t))$.

An interesting aspect of this reduction is the *time reversal*. By that we mean that if one considers a T-horizon Random-Turn Game and T steps of Synchronous Edge Balancing, then the first step of Synchronous Edge Balancing actually corresponds to the last step in the Random-Turn Game. In general, the t^{th} balancing step corresponds to t steps remaining in the game.

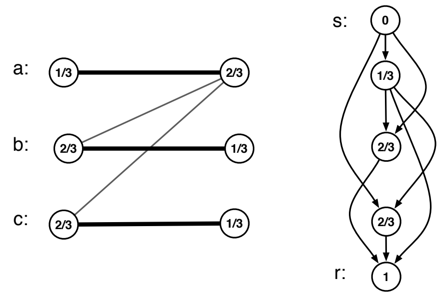

Fig. 1. An unweighted bipartite graph G and its corresponding digraph D with balanced allocations.

Throughout this paper, we say a graph D is *weakly acyclic* if the only directed cycles it contains go through s or r. If a graph G reduces to a digraph D that is weakly acyclic then the balanced outcome on G is *unique*. The converse also holds for unweighted bipartite graphs.

Consider the Random-Turn Game defined by the digraph $D = (V, E)$ and the payoff function f as above. The following theorem relates the value of the Random-Turn Game to Synchronous Edge Balancing, and shows it is sufficient to analyze the convergence of the Random-Turn Game.

Theorem 1. *The value of a* Random-Turn Game *with starting vertex v and horizon T is exactly $B(v, T)$ when the directed graph is weakly acyclic.*

Let the balanced outcome be denoted by $B(v)$. For such games, we give the *optimal* rate of convergence, which is as follows. Let h be the maximum length of a path from s to r in D.

Theorem 2. *There exists a $T \in O(h^2 \log(1/\varepsilon))$ such that for all $t \geq T$ the value of the* Random-Turn Game *starting at vertex v with horizon t is within ε of $B(v)$, given that D is weakly acyclic.*

The proof of this theorem is the most technical part of the paper, and uses techniques from the theory of martingales. Recall that an allocation f' is *ε-close to balanced* if there exists a balanced outcome B' such that $|B'(v) - f'(v)| \leq \varepsilon$ for all v. We can now restate the result and the corresponding rate of convergence in Synchronous Edge Balancing. The proofs are the focus of Section 3.

Theorem 3. Synchronous Edge Balancing *on unweighted bipartite graphs with a unique balanced outcome results in an allocation that is ε-close to a balanced outcome after at most $O(|M|^2 \log(1/\varepsilon))$ rounds of the balancing process.*

This result follows directly from Theorem 2 and the fact that $h \leq |V| = |M| + 2$. Lastly, we show our result is tight.

Theorem 4. *There exist graphs G with matchings M and initial allocations such that the balancing process requires $\Omega(|M|^2 \log(1/\varepsilon))$ time to be ε-close to a balanced outcome.*

Sketch of Convergence Proof

We now give a brief sketch of the proof of Theorem 2 for the case when M is a *perfect matching* in G. Observe that if a game with finite horizon ends in an absorbing state, then the vertex payoffs don't matter. Thus one approach is to show that with high probability, a Random-Turn Game with a sufficiently large horizon ends in an absorbing state. To be precise, let $\{X_t\}$ be a sequence of vertices in a run of the Random-Turn Game under optimal play. We wish to show that for a game with sufficiently large horizon T, $X_T \in \{s, r\}$ with high probability. However, it is unclear how to analyze the behavior of X_t. Instead we show it is sufficient to analyze the related sequence of vertices $\{Y_t\}$ obtained when Max plays optimally, but $Mini$ plays as if the payoff function was B. We show $B(Y_t)$, the value of the balanced outcome of vertex Y_t, is a supermartingale. Moreover, we know that it is bounded in $[0, 1]$ and show that its conditional variance is at least $1/h^2$. These suffice to prove the desired bound on the absorption time.

Extensions

To summarize, the approach outlined to prove convergence of Synchronous Edge Balancing is as follows: reduce it to convergence of a Random-Turn Game (Theorem 1) and show bounds on this game (Theorem 2). The first part of this approach can be extended naturally to show convergence (but not rates) for may variants of the dynamics and general graphs. For *non-bipartite graphs* we maintain a vertex in D for each matched vertex in G. If the graph is *weighted* we use running payoffs in the random-turn game. *Damped dynamics* correspond to lazy random-turn games. And if we wish vertices to be *individual rational*, then the corresponding *capped dynamics* are captured by a random-turn game where the players are allowed to quit. This list is far from exhaustive, but illustrates the flexibility and robustness of our technique and is discussed further in the full version of this paper.

3 Rate of Convergence

We begin with the proof of Theorem 1. We recall some notation: given an allocation f, $v_f^+ = \arg\max_{u:uv\in E}\{f(u)\}$ and $v_f^- = \arg\min_{u:vu\in E}\{f(u)\}$. The allocation in round t of Synchronous Edge Balancing is $B_f(v, t)$ (we now drop the subscript f for convenience). The updates are, $B(v, t+1) = \frac{1}{2}(B(v^+, t) + B(v^-, t))$ where v^+ and v^- are defined with respect to $B(v, t)$. Theorem 1 says that $B(v, T)$ is the value of the Random-Turn Game starting at vertex v with horizon T. The proof is by induction on T. We first strengthen the inductive hypothesis to assume the *optimal strategies* for Max and $Mini$ are to choose v^+ and v^- respectively. We refer to this strategy as the *balancing strategy.*

Theorem 5. *Given a* Random-Turn Game *with horizon* T*, the optimal strategy for either player is the balancing strategy.*

Proof (Theorems 1 and 5). The proof is by a joint induction on the horizon t to prove (a) $B(v,t)$ is the value of the game and (b) the optimal strategy when $t+1$ moves remain is the balancing strategy.

In the base case, $t = 0$. To show (a), note that the expected payoff of the game for Max at node v is exactly $B_f(v,0) = f(v)$ since there are no moves to be made. To show (b), consider the horizon $t+1 = 1$ at a given node v. In this case, optimal moves for Max and $Mini$ are clearly v_f^+ and v_f^- respectively, since the payoff at the end of this turn will be the terminal payoff of the game.

For the inductive step, let us assume that for all v and some $t \in \mathbb{N}$, the value of the game of horizon $t-1$ is $B_f(v,t-1)$, and in the t horizon game the bargaining strategy is optimal. To prove (a) we note that the latter statement implies Max will move to $v^+_{B(v,t-1)}$ if he wins the coin toss and $Mini$ will move to $v^-_{B(v,t-1)}$ if she wins the coin toss. From the first part of the inductive hypothesis we know $B(v,t-1)$ is the expected payoff for Max in the $t-1$ horizon game. Thus, the expected payoff of the game for Max under optimal play in the t horizon game is $\frac{1}{2}(B(v^+,t-1)+B(v^-,t-1)) = B(v,t)$. To prove (b), consider the $t+1$ horizon game. Under optimal play, Max wishes to maximize his expected payoff, and $Mini$ wishes to minimize the expected amount she has to pay. Assume we are at vertex u, and recall that Max must move to a predecessor of u and $Mini$ to a successor. Since there are t steps remaining after the initial step, an optimal strategy for Max ($Mini$) will maximize (minimize) the expected payoff $B_f(v,t)$. Thus, if Max wins the toss he will move to $v^+_{B(v,t)}$ and if $Mini$ wins it she will move to $v^-_{B(v,t)}$, which is precisely the balancing strategy. □

We now give the proof of Theorem 2 for the case where we have a perfect matching. Note that with the assumptions of the theorem, this implies D is *strongly acyclic*; i.e. it does not contain cycles of any kind. We briefly explain the technical extension for non-perfect matchings at the end of this section. The main idea behind the proof is to first reduce the analysis to showing that a particular sequence $\{Y_t\}$ (of vertices in V) gets absorbed at $\{s,r\}$ with high probability, and then show this happens in polynomial time using techniques from the theory of martingales.

Proof (Theorem 2). Consider two allocations, f and g such that $f(v) \leq g(v)$ for all v. We show in Lemma 1 that $B_f(v,t) \leq B_g(v,t)$ for all v,t. Hence, if we consider the initial allocations

$$\mathbf{0}(v) = \begin{cases} 0 \text{ if } v \neq r; \\ 1 \text{ otherwise.} \end{cases} \qquad \text{and} \qquad \mathbf{1}(v) = \begin{cases} 1 \text{ if } v \neq s; \\ 0 \text{ otherwise.} \end{cases},$$

we have $B_{\mathbf{0}}(v,t) \leq B_f(v,t) \leq B_{\mathbf{1}}(v,t)$ for all v,t, and f. Thus, it suffices to prove that $B_{\mathbf{0}}(v,T) \geq B(v) - \varepsilon$ and $B_{\mathbf{1}}(v,T) \leq B(v) + \varepsilon$ for $T \in O(h^2 \log(1/\varepsilon))$. We will prove the latter, and the proof for the former follows exactly with the roles of $Mini$ and Max reversed and the payoff function $\mathbf{0}$ instead of $\mathbf{1}$.

Consider the game with payoff function $\mathbf{1}$ and horizon T where $T \in O(h^2 \log(1/\varepsilon))$. Consider the sequence of vertices $\{X_t\}$ with $X_0 = v$ that occurs if $Mini$ and

Max play optimally. From Theorem 1,

$$B_{\mathbf{1}}(v,t) = \mathbb{E}\mathbf{1}(X_t). \tag{1}$$

Now consider the half-optimal sequence $\{Y_t\}$ with $Y_0 = v$, where Max plays optimally for the payoff function **1** and $Mini$ plays optimally for the payoff function B. For the game with payoffs **1** Max's expected payoff is only higher. That is

$$\mathbb{E}\mathbf{1}(X_t) \leq \mathbb{E}\mathbf{1}(Y_t). \tag{2}$$

Our key result in Lemma 3 shows that for any function f, $\mathbb{E}_v[|f(Y_T) - B(Y_T)|] \leq \varepsilon$. (The proof of this lemma follows by showing convergence of the sequence $\{Y_t\}$.) If we take $f = \mathbf{1}$ and note that $\mathbf{1}(Y_t) \geq B(Y_t)$, we get

$$\mathbb{E}\mathbf{1}(Y_T) \leq \mathbb{E}B(Y_T) + \varepsilon. \tag{3}$$

Now consider the sequence $\{Z_t\}$ with $Z_0 = v$ that occurs when $Mini$ and Max play optimally for the payoff function B. The expected payoff for Max with payoff function B is higher in $\{Z_t\}$ than in $\{Y_t\}$. Thus

$$\mathbb{E}B(Y_T) \leq \mathbb{E}B(Z_T). \tag{4}$$

Finally, we show in Lemma 2 that

$$\mathbb{E}B(Z_T) = B(v). \tag{5}$$

From (1) – (5), it follows that $B_{\mathbf{1}}(v,T) \leq B(v) + \varepsilon$ as desired. □

Lemma 1. *The balancing process is monotonic, namely if $f(v) \leq g(v)$ for all $v \in V$, then $B_f(v,t) \leq B_g(v,t)$ for all v, t.*

Lemma 2. *The value of a* Random-Turn Game *with function $f = B$ is equal to B for all horizons $T \in \mathbb{N}$.*

This Lemma follows from Theorem 1 and the observation that B is a fixed point of Synchronous Edge Balancing. A detailed proof of both lemmas can be found in the full version of this paper.

Lemma 3. *Consider the expected payoff for Max in the half-optimal chain $\{Y_t\}$ defined in the proof of Theorem 2. For sufficiently large t, the expected payoff for Max with payoff function f is close to the balanced outcome B. Specifically, $\mathbb{E}_v[|f(Y_T) - B(Y_T)|] \leq \varepsilon$ when $T \geq 4h^2 \log(1/\varepsilon)$.*

Proof. Clearly if $Y_t \in \{s, r\}$, then the game has ended and $f(Y_t) - B(Y_t) = 0$. Additionally, the difference $|f(Y_t) - B(Y_t)|$ is at most 1 since $f(v), B(v) \in [0,1]$ for all $v \in V$. Thus, the expected difference $\mathbb{E}_v[|f(Y_t) - B(Y_t)|]$ is at most the probability that Y_t has not been absorbed.

Let us now show this probability is bounded by ϵ, namely $\mathbb{P}r_v[Y_t \notin \{s,r\}$ for $t \geq 4h^2 \log(1/\varepsilon)] \leq \varepsilon$ for all $v \in V$. The main convergence is shown in Lemma 4 which says that the probability that $Y_t \notin \{s,r\}$ for $t = 4h^2$ is at most $\frac{1}{4}$. Since the statement holds for all $v \in V$, if we are not at s or r after $4h^2$ time steps we can simply apply the lemma again. Thus, after $4h^2 \log(1/\varepsilon)$ time steps, the probability that we are not at s or r is $(\frac{1}{4})^{\log(1/\varepsilon)} = 4^{\log \varepsilon} \leq \varepsilon$. □

Lemma 4. *$\mathbb{P}r_v[Y_t \notin \{s, r\}$ for $t \geq 4h^2] \leq \frac{1}{4}$ for all $v \in V$ where h is the height of D[7] and $\{Y_t\}$ is the half-optimal chain defined above.*

Proof. Let the absorption time be $\tau = \min\{t : Y_t \in \{s, r\}\}$. Note that $\mathbb{P}r[Y_t \notin \{s, r\}$ for some $t \geq 4h^2] = \mathbb{P}r[\tau \geq 4h^2]$. We show that $\mathbb{E}[\tau] \leq h^2$. Then by Markov's inequality, $\mathbb{P}r[\tau \geq 4h^2] \leq \frac{1}{4}$ as desired.

Consider the sequence $\{\Psi_t\} = \{B(Y_t)\}$. In the half-optimal chain $\{Y_t\}$, Max plays suboptimally and $Mini$ plays optimally according to payoff function B (see the proof of Theorem 2). Hence $B(Y_t)$ is an upper bound on the expected payoff for Max at time t, and therefore $\{\Psi_t\}$ is a supermartingale.[8]

Now consider the quadratic chain $\Phi_t = 2\Psi_t - \Psi_t^2 + t\sigma^2$ where σ^2 is a lower bound on the conditional variance of Ψ_t. We show that Φ_t is also a supermartingale (Lemma 5). Therefore, since $\Phi_t \geq 0$, the optional stopping theorem[9] gives $\mathbb{E}[\Phi_\tau] \leq \Phi_0 \leq 1$. The bounds on Ψ_t also imply that $2\Psi_t - \Psi_t^2 \geq 0$, and hence we get $\mathbb{E}[\Phi_\tau] \geq \mathbb{E}[\tau]\sigma^2$,and $\mathbb{E}[\tau] \leq \frac{1}{\sigma^2}$. By Lemma 6, we know that we can take $\sigma^2 = \frac{1}{h^2}$, so $\mathbb{E}[\tau] \leq h^2$ as required. □

Lemma 5. *Given a supermartingale $0 \leq \Psi_t \leq 1$ with conditional variance at least σ^2, the quadratic chain $\Phi_t = 2\Psi_t - \Psi_t^2 + t\sigma^2$ is a supermartingale.*

Lemma 6. *The variance of a step in $\{\Psi_t\}$ is at least $\sigma^2 = 1/h^2$.*

The proofs of these lemmas can be found in the full version of this paper.

When M is not a perfect matching, D is a weakly acyclic digraph with a cycles through s and/or r. Any vertex in such a cycle must take value exactly 0 or 1 in the balanced outcome, thus these cycles can be treated as absorbing states. Hence, we can first analyze the mixing time of the cycle using spectral techniques[10], and then apply the theorems above to get the same time bound.

4 Conclusion and Future Work

We reduced the problem of analyzing the convergence of local dynamics for a network bargaining game to that of a random-turn game. With this reduction we bring all the machinery from the analysis of random processes, especially the theory of Markov chains and martingales, to the analysis of local dynamics. We used these techniques to give the optimal bound on unweighted graphs with a unique balanced outcome. Prior to this work, there was no effective technique known to analyze such dynamics, and the best bound known on any non-trivial class of graphs was exponential.

Our work opens up a promising line of approach to analyze many variants of local dynamics on general graphs. The most immediate is perhaps to bound the convergence rate for weighted graphs. The difficulty with our current analysis is

[7] The *height* is the length of the longest path from s to r.
[8] Recall that a supermartingale is a sequence $\{a_t\}$ in which $a_t \geq \mathbb{E}[a_{t+1}|a_t]$.
[9] See Theorem 10.10 (d) in [36].
[10] See Chapter 12 in Peres, et al. [26] for an exposition on spectral techniques.

that the supermartingale Ψ_t we used in the unweighted case is unbounded when there are weights. We believe a different supermartingale that does not suffer from this drawback could give the appropriate bound.

The most significant technical hurdle arises when D is cyclic. In this case, the game may never end, since it might get stuck in a *stalemate*, where the players travel in a cycle indefinitely. Thus, a bound on the absorption time of the game does not suffice – we must analyze the behavior on the cycle separately by internally considering its *mixing time*, and externally treating it as an absorbing state.[11] However, the details of such an analysis remain unclear.

A final important direction is to obtain tight polynomial bounds for dynamics which find both the matching and the balanced outcome simultaneously. One approach would be to combine the dynamics by Kanoria et. al. [22] with our techniques to attain a tight polynomial rate of convergence.

References

1. Workshop on new connections between differential and random turn games, pde's and image processing (2009)
2. Ackermann, H., Goldberg, P.W., Mirrokni, V.S., Röglin, H., Vöcking, B.: Uncoordinated two-sided matching markets. In: EC. pp. 256–263 (2008)
3. Azar, Y., Birnbaum, B., Celis, L.E., Devanur, N.R., Peres, Y.: Convergence of local dynamics to balanced outcomes in exchange networks. In: FOCS (2009)
4. Azar, Y., Devanur, N., Jain, K., Rabani, Y.: Monotonicity in bargaining networks. In: SODA (2010)
5. Bateni, M.H., Hajiaghayi, M.T., Immorlica, N., Mahini, H.: Network bargaining with capacity constraints. In: ICALP (2010)
6. Bayati, M., Borgs, C., Chayes, J.T., Zecchina, R.: Belief-propagation for weighted b-matchings on arbitrary graphs and its relation to linear programs with integer solutions. SIAM J. in Discrete Math, to appear (2010)
7. Berlekamp, E.R., Conway, J.H., Guy, R.K.: Winning ways for your mathematical Plays. Academic Press, London (1982)
8. Chakraborty, T., Judd, S., Kearns, M., Tan, J.: A behavioral study of bargaining in social networks. In: EC (2010)
9. Chakraborty, T., Kearns, M.: Bargaining solutions social networks. In: WINE (2008)
10. Chakraborty, T., Kearns, M., Khanna, S.: Nash bargaining: Algorithms and structural results. In: EC (2009)
11. Charro, F., García Azorero, J., Rossi, J.D.: A mixed problem for the infinity Laplacian via tug-of-war games. Var. & Par. Differential Eq. 34(3), 307–320 (2009)
12. Chen, X., Deng, X.: Settling the complexity of two-player Nash equilibrium. In: FOCS. pp. 261–272 (2006)
13. Cook, K.S., Emerson, R.M., Gillmore, M.R., Yamagishi, T.: The distribution of power in exchange networks: Theory and experimental results. American Journal of Sociology 89, 275–305 (1983)
14. Cook, K.S., Yamagishi, T.: Power in exchange networks: a power-dependence formulation. Social Networks 14(3-4), 245–265 (1992)

[11] This proposed approach is a generalization of our analysis of non-perfect matchings in unweighted graphs.

15. Corominas-Bosch, M.: Bargaining in a network of buyers and sellers. Journal of Economic Theory 115(1), 35 – 77 (2004)
16. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. In: STOC. pp. 71–78 (2006)
17. Draief, M., Vojnovic, M.: Bargaining dynamics in exchange networks. In: The 48th Annual Allerton Conference on Communication, Control, and Computation. Monticello, IL (September 2010)
18. Faigle, U., Kern, W., Kuipers, J.: An efficient algorithm for nucleolus and prekernel computation in some classes of TU-games. Ph.D. thesis, Memorandum No. 1464, University of Twente (1998)
19. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. The American Mathematical Monthly 69(1), 9–15 (1962)
20. Jain, K.: A polynomial time algorithm for computing an Arrow-Debreu market equilibrium for linear utilities. SIAM J. Comput. 37(1), 303–318 (2007)
21. Kanoria, Y.: An FPTAS for bargaining networks with unequal bargaining powers. In: WINE, to appear (2010)
22. Kanoria, Y., Bayati, M., Borgs, C., Chayes, J.T., Montanari, A.: Fast convergence of natural bargaining dynamics on exchange networks. In: SODA, to appear (2011)
23. Kleinberg, J., Tardos, E.: Balanced outcomes in social exhchange networks. In: STOC (2008)
24. Lazarus, A., Loeb, D.E., Propp, J.G., Stromquist, W.R., Ullman, D.H.: Combinatorial games under auction play. Games and Economic Behavior 27 (1999)
25. Lazarus, A., Loeb, D.E., Propp, J.G., Ullman, D.H.: Richman games. Games of No Chance 29, 439–449 (1996)
26. Levin, D.A., Peres, Y., Wilmer, E.L.: Markov Chain Mixing Times. Academic Mathematical Society, London (2009)
27. Maschler, M.: The bargaining set, kernel, and nucleolus. In: Aumann, R., Hart, S. (eds.) Handbook of Game Theory with Economic Applications, vol. 1, chap. 18, pp. 591–667. Elsevier (July 1992)
28. Nash, J.: The bargaining problem. Econometrica 18, 155–162 (1950)
29. Paulusma, D.: Complexity Aspects of Cooperative Games. Ph.D. thesis, University of Twente, Enschede, the Netherlands (2001)
30. Peres, Y., Schramm, O., Sheffield, S., Wilson, D.B.: Tug-of-war and the infinity laplacian. Journal of the AMS 22, 167–210 (2009)
31. Rochford, S.C.: Symmetrically pairwise-bargained allocations in an assignment market. Journal of Economic Theory 34(2), 262–281 (1984)
32. Sanghavi, S., Malioutov, D.M., Willsky, A.S.: Linear programming analysis of loopy belief propagation for weighted matching. In: NIPS (2007)
33. Shapley, L.S., Shubik, M.: The assignment game I: The core. Int. Journal of Game Theory 1(2), 111–130 (1972)
34. Solymosi, T., Raghavan, T.E.S.: An algorithm for finding the nucleolus of assignment games. Int. Journal of Game Theory 23(2), 119–143 (1994)
35. Willer, D. (ed.): Network Exchange Theory. Praeger (1999)
36. Williams, D.: Probability with Martingales. Cambridge University Press (1991)
37. Yu, Y.: Maximal and minimal solutions of an Aronsson equation: L^∞ variational problems versus the game theory. Var. & Par. Differential Eq. 37, 63–74 (2010)