Local Two-Stage Myopic Dynamics for Network Formation Games

Esteban Arcaute[1], Ramesh Johari[2], and Shie Mannor[3]

[1] Institute for Computational and Mathematical Engineering, Stanford University
arcaute@stanford.edu
[2] Department of Management Science and Engineering, Stanford University
rjohari@stanford.edu
[3] Department of Electrical and Computer Engineering, McGill University
shie@ece.mcgill.ca

Abstract. *Network formation games* capture two conflicting objectives of self-interested nodes in a network. On one hand, such a node wishes to be able to reach all other nodes in the network; on the other hand, it wishes to minimize its cost of participation. We focus on myopic dynamics in a class of such games inspired by transportation and communication models. A key property of the dynamics we study is that they are *local*: nodes can only deviate to form links with others in a restricted neighborhood. Despite this locality, we find that our dynamics converge to efficient or nearly efficient outcomes in a range of settings of interest.

1 Introduction

Viewing modern data-networks, such as the Internet or ad-hoc networks, as a federation of selfish, independent actors leads to a range of interesting game theoretic questions. The goal of a selfish agent in a data network is two-fold. On one hand, it wishes to be able to reach all other agents in the network. On the other hand, it wishes to minimize its cost for participating in the network. The recent literature on *network formation games* (NFGs) provides a natural model with which to study this tradeoff between cost and connectivity; NFGs have been suggested as models for many domains, from trade networks to peer-to-peer networks (see [1] for a comprehensive review).

In the model we consider, nodes derive utility from connectivity to each other, and incur a cost comprised of three terms: (1) traffic related costs; (2) costs to maintain links to other nodes; and (3) payments made to other nodes. The payments are a natural mechanism for users to compensate each other for the costs required to establish links; we assume these payments are bilaterally negotiated through *contracts*. Because link formation is bilateral in NFGs, the equilibrium concept of interest for NFGs is *pairwise Nash stable equilibrium* [2], as opposed to Nash equilibrium in traditional game theoretic models. Roughly speaking, a graph is pairwise Nash stable if it is a Nash equilibrium, and if it is *pairwise stable* as first introduced in [3]: no node can profitably remove edges adjacent to it, and no two nodes can deviate jointly by altering their contract to improve both their payoffs.

This paper studies dynamic processes of network formation in the setting described above; in particular, we study whether realistic dynamics exist that also naturally select

C. Papadimitriou and S. Zhang (Eds.): WINE 2008, LNCS 5385, pp. 263–277, 2008.
© Springer-Verlag Berlin Heidelberg 2008

efficient equilibria. The most basic dynamic network formation process is *best response dynamics*, where at each round a profitable deviation is undertaken by one or a pair of nodes at a time. An outcome is a fixed point of best response dynamics if and only if it is a pairwise stable equilibrium. Although best response dynamics are attractively simple, they may fail to converge; further, since any pairwise Nash stable equilibrium is a fixed point, best response dynamics can lead to inefficient equilibria.

Our main insights are that, in a model we previously introduced in [4], there exist a natural class of dynamics that also select efficient equilibria. When there is no cost to maintain links and some equilibria with redundant edges exist, our dynamics can select those non-tree equilibria; however, in this case the resulting equilibria may be inefficient.

The dynamics we study are a generalization of local-best response dynamics. First, nodes only deviate with other nodes in a local neighborhood; this is consistent with the observation that in modern data networks, nodes are typically only aware of the local topology of the network around them. Second, we allow nodes to deviate *twice* in succession, using the first deviation to improve its bargaining power in the second step.

Formally, the dynamics proceed in *rounds*. Each round is divided into two stages. During the first stage, an exogenously designated node u considers all possible unilateral deviations. During the second stage, the same node considers all possible unilateral and bilateral deviations with nodes in its local neighborhood. The key assumption is that u seeks to maximize its prospective payoff *at the end of the round*, while nodes involved in bilateral deviations choose their actions in order to maximize prospective payoff *at the end of the stage*, considering the current network when making a decision. As such, one can think of our dynamics as constrained local-best response dynamics with *one-step look-ahead* for the active node. It is straightforward to show in the models we consider that any fixed point of these dynamics must be pairwise Nash stable.

The one-step look-ahead feature allows the deviating node more flexibility than best-response dynamics—and thus in principle encourage more selfish behavior. The main benefit of the proposed dynamics is that they serve as a local and decentralized equilibrium selection mechanism. In particular, we show that for the utility model from [4], where potentially inefficient pairwise Nash stable equilibria exist, these dynamics in fact converge to desirable pairwise Nash stable equilibria instead. Also, when there is no link maintenance cost and equilibria with redundant edges exist, we prove that it is possible for the most inefficient equilibria selected to be less efficient than those selected when there is a positive link-maintenance cost. Thus we prove that there is a possible efficiency loss when the link maintenance cost vanishes.

In the NFG model we consider, a node incurs a traffic related cost that depends on the total volume of traffic routed through it. In [4], it was shown that, for a very restricted class of global dynamics, the network converges to a desirable pairwise Nash stable equilibrium under certain conditions. Our paper recovers this convergence result and extends it to the case when the link maintenance cost is nil using significantly more general and natural, local dynamics; at the same time, the proofs are simplified via an ordinal potential function argument (see [2]).

Related Work

Our work touches on several related threads of the literature. Most closely related is the work on network formation games in economics (see [1] for a survey). In the context of communication networks [5] considered a static network formation game related to the model in this paper; by contrast, the focus of the current paper is on the dynamics of such formation processes. The work of Jackson and Watts also considers dynamics for network formation games [6, 7], but for a utility model that is unrelated to ours. While in their dynamics only a unilateral or a bilateral deviation may occur in a given round, our dynamics are designed so that each round consists of two stages, thus allowing a unilateral deviation to be *followed* by a bilateral deviation. This latter property allows our dynamics to *select* desirable equilibria. Despite the dissimilarities in cost structure and dynamics, random activation of nodes in the dynamics is needed for obtaining convergence results both in our setting and in that of [7].

In [4], we considered essentially the same utility model and static game as in this paper. There are several important differences that we summarize bellow.

– In [4] we used a variant of pairwise stability as the solution concept. Here we use pairwise Nash stability, which allows us to define our dynamics in a more natural framework, and also allows us to provide a complete characterization of equilibrium outcomes.
– In [4] we considered a restricted version of the dynamics we introduce in this paper. Each round is again composed of two stages. In [4], at the beginning of each round, an edge is sampled from the set of all possible edges, the active node is selected among the two end-points of the sampled edge. Two main differences can be pointed out:
    - During the first stage, the active node is allowed to unilaterally deviate *only* with respect to the sampled edge. This is in sharp contrast to the dynamics of this paper where the active node can unilaterally deviate with respect to *any* set of edges adjacent to it.
    - During the second stage, the active node can bilaterally deviate with a node selected from *all* possible nodes, thus requiring all nodes to be aware of the global topology of the network. In our current dynamics, the active node can choose to unilaterally or bilaterally deviate, and can only do so with nodes in its local neighborhood.

  Not only are the dynamics in [4] a very restricted version of those considered in this paper (as the type of deviations allowed during each stage are explicitly restricted), but also the nodes are required to have *global information* about the network (which is not realistic in most modern data-networks).

  The main result of [4] is that the network converges to a stable and desirable equilibrium under certain conditions; we recover this result here, but for the more general model of dynamics.
– Further, in this paper we are able to prove a convergence result even in the setting where the link formation cost may be zero. This result is nontrivial because, as noted, the limiting topologies may not be trees (as is always the case in the analysis of [4]).

In contrast to our approach, which is inspired by the literature on network formation games, in recent years a large body of literature has considered efficiency loss due to self-interested behavior in routing-related applications; see, e.g., [8]. A related thread of literature has considered the efficiency of equilibria in *network design games*; see, e.g., [9, 10, 11]. In this line of work, a network is formed based on *unilateral* decisions of the players, and the cost of the resulting network is shared among the players. The goal of that body of research is to design mechanisms that incentivize agents to choose to use the network in an efficient manner.

Our work is also related to the literature on learning in games; see, e.g., [12, 13, 14] for surveys. In this literature the emphasis is on studying classes of dynamic methods that converge into the set of equilibria (e.g., correlated or Nash equilibria), without regard to efficiency. Our approach departs significantly from this literature, as we are interested in convergence to desirable equilibria.

Finally, there is an extensive body of research in the application of game theory to networks; see, e.g., [15] for a survey, and [16, 17] for a discussion of pricing in networks. In the application domain, our work is related to papers on topology formation in ad hoc networks; e.g., [18, 19, 20]. However, these works all consider Nash equilibria, whereas our focus is on *pairwise* interactions between nodes.

The rest of the paper is organized as follows. In Section 2 we present the notation used in the paper. In Section 3 we present the utility model considered In Section 4 we define the static game, and recall the definition of pairwise Nash stability. In Section 6 we define the proposed two-stage dynamics. The main results are stated in Section 7 and are discussed in Section 8. The proofs of the main results are in online appendices in [21].

2 Notation

Let $G(V, E)$ be an undirected graph over V, with edge set E. We assume that V has n elements, and E has m elements. For a given pair of nodes u and v in V, we call $d(u, v; G)$ the distance in G (i.e., number of edges) from u to v. If u and v are in two distinct connected components, we set the distance $d(u, v; G)$ to be n. Finally, for a given node $u \in V$, we call $\delta_G(v)$ its degree in G.

3 Utility Model

We assume the nodes of the graph are strategic agents. We assume the cost to a node $v \in V$ is of the form:

$$C(v; G) = \pi\delta_G(v) + h(v; G), \tag{1}$$

for some positive function h. We assume that the cost is arbitrarily large (but finite) if G is disconnected; thus we will restrict our model and analysis to connected networks.

The function h models a situation where the cost to a node v is proportional to $f(v; G)$, the number of packets v forwards and receives in G.

We now define more precisely the function h.

Definition 1 (Routing-Related Cost). *Let $v \in V$ be a node in a connected network G. Let $c_v > 0$ be given. We call c_v the* per-unit routing cost *associated to v. We define the routing-related cost to v in G as*

$$h(v; G) = c_v f(v; G), \tag{2}$$

where $f(v, G)$ is determined by the routing policy and the traffic matrix. In our setting, we assume all-to-all traffic routed uniformly over all shortest paths.

Next, we assume that links in the network are the result of bilateral agreements between the participating nodes. We further assume that such an agreement induces a utility transfer, or payment, between the nodes participating in the agreement. Let $\mathbf{P}$ be a n by n matrix where, for all $1 \leq i, j \leq n$, p_{ij} is the utility transfer from node i to node j. We call $\mathbf{P}$ the *payment matrix*.

Let $G(V, E)$ be a given connected network topology, and let (P) be a given n by n matrix with real entries, and $P_{ii} = 0$ for all i. The total utility $U(i; G, \mathbf{P})$ to node i in G, given the payment matrix $\mathbf{P}$, is:

$$U(i; G, \mathbf{P}) = \sum_{j \neq i} (p_{ji} - p_{ij}) - C(i; G), \tag{3}$$

i.e., the total payments made to i, minus the total payments made by i, minus the cost of being in network topology G.

3.1 Contracting

In order to completely define the total utility to node i as in Equation 3, we need to establish how entries in the payment matrix $\mathbf{P}$ are calculated. Recall that we assumed that a link $e = ij$ in G is the result of a bilateral agreement between i and j. We assume that the link arises due to a *contract* between i and j. Contracts are directional, i.e., the link ij may exist due to either the contract (i, j) or the contract (j, i).

If $ij \notin G$, and the contract (i, j) is agreed upon by both i and j, then we assume that i pays j an amount $Q(i, j; G + ij)$; the function Q here is called the *contracting function*. This function gives the value of the contract formed between i and j, given the resulting network topology.

We believe two interpretations of the contracting function are reasonable. First, we might imagine that an external regulator has dictated that contracts between nodes must have pre-negotiated tariffs associated with them; these tariffs are encoded in the contracting function. Note that the regulator in this case dictates changes in the value of the contract as the surrounding network topology changes.

A second interpretation of the contracting function does not assume the existence of the regulator; instead, we presume that the value of the contracting function is the outcome of bilateral negotiation between the nodes in the contract. Note that the structure of our game assumes that this negotiation takes place *holding the network topology fixed*; i.e., the negotiation is used to determine the value of the contract, given the topology that is in place. One example is simply that $Q(i, j; G)$ is the result of a Rubinstein bargaining game of alternating offers between i and j, where i makes the first offer

[22]. If both players are infinitely patient, the resulting contracting function is identical to uniform cost sharing. More details can be found in [4].

We will be interested in contracting functions exhibiting two natural properties: *monotonicity* and *anti-symmetry*. Informally, monotonicity states that, given a network topology, the utility transfer associated with a contract is increasing in the burden associated to the contract proposed. Anti-symmetry asserts that $Q(i, j; G)$ must be equal to the negation of $Q(j, i; G)$. Anti-symmetric contracting functions ensure that directionality of the contract does not affect the value of the contract (i.e., (i, j) and (j, i) lead to the same utility transfer).

Assumption 1 (Anti-symmetry). *The contracting function is anti-symmetric, i.e., for all $u \neq v$ and for all G,*

$$Q(u, v; G + uv) = -Q(v, u; G + uv)$$

Assumption 2 (Monotonicity). *We assume that the contracting function is monotone in the change of traffic cost to a node. In other words, let u, v and w be three distinct nodes such that $uv \notin G$ and $uw \notin G$. Let G_u, G_v and G_w be the connected components where u, v and w lie respectively. Then if $h(w; G_u \cup G_w + uw) - h(w; G_w) < h(v; G_v \cup G_u + uv) - h(v; G_v)$, then $Q(u, w; G + uw) < Q(u, v; G + uv)$; and if $h(w; G_u \cup G_w + uw) - h(w; G_w) = h(v; G_v \cup G_u + uv) - h(v; G_v)$, then $Q(u, w; G + uw) = Q(u, v; G + uv)$.*

3.2 State of the Game

From Equation 3 we see that the utility of all nodes is defined by G and $\mathbf{P}$. In order to keep track of the contracts in place at a given time in the network, we introduce the *contracting graph* Γ to be a directed graph over V such that $(i, j) \in \Gamma$ if nodes i and j have agreed to the corresponding contract.

Thus the state of the system is completely determined by the tuple $(G, \mathbf{P}, \Gamma)$.

Remark 1. Note that, from our contracting assumption, we have the following properties:

1. for all $i \neq j$, $ij \in G$ if and only if $(i, j) \in \Gamma$ or $(j, i) \in \Gamma$; and
2. for all $i \neq j$, $p_{ij} \neq 0$ implies $(i, j) \in \Gamma$.

4 Static Game

We can now formally define the game we consider. We use the same static game we defined in [4]. Note that variants of this game have been considered in previous literature (see [23, 24, 5]). We consider a network formation game in which each node selects nodes it wishes to connect to, as well as nodes it is willing to accept connections from. Formally, each node i simultaneously selects a subset $F_i \subseteq V$ of nodes i is willing to accept connections from, and a subset $T_i \subseteq V$ of nodes i wishes to connect to. We let $\mathbf{T} = (T_i, i \in V)$ and $\mathbf{F} = (F_i, i \in V)$ denote the composite strategy vectors.

An undirected link is formed between two nodes i and j if i wishes to connect to j (i.e., $j \in T_i$), and j is willing to accept a connection from i (i.e., $i \in F_j$). All edges that are formed in this way define the network topology $G(\mathbf{T}, \mathbf{F})$ realized by the strategy vectors $\mathbf{T}$ and $\mathbf{F}$; i.e., $j \in T_i$, $i \in F_j$ implies that $ij \in G(\mathbf{T}, \mathbf{F})$.

Further, if $i \in F_j$ and $j \in T_i$, then a *binding contract* is formed from i to j; we denote this contract by (i, j), and add it to the contracting graph $\Gamma(\mathbf{T}, \mathbf{F})$. The contracting graph captures the inherent directionality of link formation: in our model a link is only formed if one node asks for the link, and the target of the request accepts.

Finally, given a contracting function Q, as presented in Section 3.1, we define the payment matrix $\mathbf{P}(\mathbf{T}, \mathbf{F})$: if the formed network is G, then the contract (i, j) leads to a payment $Q(i, j; G)$ from i to j. This completely defines the state of the system. The utility of each node is defined as in Section 3, and thus the outcome of the game is well defined.

By an abuse of notation, and where clear from context, we will often use the shorthand $G = G(\mathbf{T}, \mathbf{F})$, $\Gamma = \Gamma(\mathbf{T}, \mathbf{F})$, and $\mathbf{P} = \mathbf{P}(\mathbf{T}, \mathbf{F})$ to represent specific instantiations of the network topology, contracting graph, and payment matrix, respectively, arising from strategy vectors $\mathbf{T}$ and $\mathbf{F}$. We refer to a triple $(G, \Gamma, \mathbf{P})$ arising from strategic decisions of the nodes as a *feasible outcome* if there are strategy vectors $\mathbf{T}$ and $\mathbf{F}$ that give rise to $(G, \Gamma, \mathbf{P})$.

5 Stability and Efficiency

In this section we define our solution concept, *pairwise Nash stability*, and formally define efficiency. While this is related to our earlier work in [4], the modification of our solution concept is an important change from our previous work: it allows us to present the dynamics in a succinct way (as will be clear in Section 6), and to better interpret the results (as will be clear in Section 8).

As is commonly observed in network formation games, Nash equilibrium lacks predictive power because link formation is inherently a *bilateral* process; thus we adopt the notion of *pairwise Nash stability* as our solution concept [2]. Informally, pairwise Nash stability requires that no unilateral deletion of contracts by a single node are profitable, and that no two nodes can simultaneously increase their utility by adding new contract(s) between them. In that sense, a network is pairwise Nash stable if it is a Nash network and pairwise stable (as originally defined by Jackson and Wolinsky in [3]).

Formally, suppose that the current strategy vectors are $\mathbf{T}$ and $\mathbf{F}$, and the current network topology and contract graph are $G = G(\mathbf{T}, \mathbf{F})$ and $\Gamma = \Gamma(\mathbf{T}, \mathbf{F})$ respectively. First, suppose that node i attempts to unilaterally deviate. Then the strategy (T_i', F_i') if i after deviation is such that $(T_i', F_i') \subseteq (T_i, F_i)$. [1] Next, suppose that two nodes i and j attempt to bilaterally deviate; this involves changing the pair of strategies (T_i, F_i) and (T_j, F_j) together such that, after the deviation, $j \in T_i'$ & $i \in F_j'$ or $i \in T_j'$ & $j \in F_i'$. Any deviation will of course change both the network topology, as well as the contract graph.

[1] Although we restrict a unilateral deviation to only encompass the deletion of contracts, it can be shown that this is equivalent to allowing *any* unilateral deviation.

However, we assume that any contracts present both before and after the deviation *retain the same payment*. This is consistent with the notion of a contract: unless the deviation by i and j entails either breaking an existing contract or forming a new contract, there is no reason that the payment associated to a contract should change. With this caveat in mind, we formalize our definition of pairwise Nash stability as follows.

Definition 2. *Assume Q is a contracting function. Given strategy vectors $\mathbf{T}$ and $\mathbf{F}$, let $G = G(\mathbf{T}, \mathbf{F})$, $\Gamma = \Gamma(\mathbf{T}, \mathbf{F})$, and $\mathbf{P} = \mathbf{P}(\mathbf{T}, \mathbf{F})$. Given strategy vectors $\mathbf{T}'$ and $\mathbf{F}'$, define $G' = G(\mathbf{T}', \mathbf{F}')$ and $\Gamma' = \Gamma(\mathbf{T}', \mathbf{F}')$. Define $\mathbf{P}'$ according to:*

$$P'_{k\ell} = \begin{cases} P_{k\ell}, & \text{if } (k,\ell) \in \Gamma' \text{ and } (k,\ell) \in \Gamma; \\ Q(k,\ell;G'), & \text{if } (k,\ell) \in \Gamma' \text{ and } (k,\ell) \notin \Gamma; \\ 0, & \text{otherwise}. \end{cases} \tag{4}$$

Then $(\mathbf{T}, \mathbf{F})$ is a pairwise Nash stable equilibrium *if: (1) No unilateral deviation is profitable, i.e., for all i, and for all $\mathbf{T}' \subseteq \mathbf{T}$ and $\mathbf{F}' \subseteq \mathbf{F}$ that differ from $\mathbf{T}$ and $\mathbf{F}$ (respectively) only in the i'th components,*

$$U_i(\mathbf{P}, G) \geq U_i(\mathbf{P}', G');$$

and (2) no bilateral deviation is profitable, i.e., for all pairs i and j, and for all $\mathbf{T}'$ and $\mathbf{F}'$ that differ from $\mathbf{T} \subseteq \mathbf{T}'$ and $\mathbf{F} \subseteq \mathbf{F}'$ only in the i'th and j'th components,

$$U_i(\mathbf{P}, G) < U_i(\mathbf{P}', G') \implies U_j(\mathbf{P}, G) > U_j(\mathbf{P}', G').$$

Notice that (4) is a formalization of the discussion above. The first condition in the definition ensures no unilateral deviation is profitable, and the second condition ensures that if node i benefits from a bilateral deviation with j, then node j must be strictly worse off.

We will typically be interested in pairwise Nash stability of the network topology and contracting graph, rather than pairwise Nash stability of strategy vectors. We will thus say that a feasible outcome $(G, \Gamma, \mathbf{P})$ is a *pairwise Nash stable outcome* if there exists a pair of strategy vectors $\mathbf{T}$ and $\mathbf{F}$ such that (1) $(\mathbf{T}, \mathbf{F})$ is a pairwise Nash stable equilibrium; and (2) $(\mathbf{T}, \mathbf{F})$ give rise to $(G, \Gamma, \mathbf{P})$. Note that by our definition of the game, for all i and j such that $(i,j) \in \Gamma$ we must have $P_{ij} = Q(i,j;G)$ in a pairwise Nash stable outcome.

We are also interested in system-wide performance from a global perspective, and for this purpose we must study the *efficiency* of pairwise Nash stable equilibria; we measure the efficiency of a network topology via the total value obtained by all nodes using that topology.

Given two feasible outcomes $(G, \Gamma, \mathbf{P})$ and $(G', \Gamma', \mathbf{P}')$, we say that $(G, \Gamma, \mathbf{P})$ Pareto dominates $(G', \Gamma', \mathbf{P}')$ if all players are better off in $(G, \Gamma, \mathbf{P})$ than in $(G', \Gamma', \mathbf{P}')$, and at least one is strictly better off. A feasible outcome is *Pareto efficient* if it is not Pareto dominated by any other feasible outcome. Since payoffs to nodes are *quasilinear* in our model, i.e., utility is measured in monetary units, it is not hard to show that a feasible outcome $(G, \Gamma, \mathbf{P})$ is Pareto efficient if $G \in \arg\min_{G'} S(G')$, where $S(G)$ is the *social cost function*:

$$S(G) = \sum_{i \in V} C_i(G).$$

We call such feasible outcomes *efficient*. (Note that, in particular, the preceding condition does not involve the contracting function; contracts induce zero-sum monetary transfers among nodes, and do not affect global efficiency.)

Let $V_0 = \{u \in V : \forall v \in V,\ c_u \leq c_v\}$, and for $u \in V_0$, let $c_{\min} = c_u$. In [25] we proved that, for $\pi > c_{\min}$, all efficient outcomes were stars centered around nodes in V_0. Thus, in such settings, all efficient outcomes have the same number of edges.

An immediate consequence from the definition of pairwise Nash stability is that a node has to either delete edges or add an edge during a deviation, but not both. By assumption, all nodes experience an arbitrarily large cost when in a disconnected network. Thus we can prove the following important theorem.

Theorem 1 (Pairwise Nash Stable Networks). *Assume that G is connected, and assume $\pi > 0$. For any contracting function Q, there exists a pairwise Nash stable outcome $(G, \Gamma, \mathbf{P})$ if and only if G is a tree.*

Hence all pairwise Nash stable networks also have the same number of edges. Thus, whenever $\pi > c_{\min}$, we define the *efficiency ratio* of a given tree T as the ratio $S'(T)/S'(G_{\text{eff}})$ where S' is equal to $S - 2(n-1)\pi$ (i.e. $S'(T)$ is the cost of routing traffic through T), and G_{eff} is the network topology in an efficient outcome.

6 Dynamics

Although the utility model discussed in Section 5 is essentially the same as that from [4] (with minor modifications), the dynamics considered in this paper constitute a major departure from our previous work in [4]. We describe our new dynamics in this section.

Before we begin, note a direct consequence of Theorem 1 is that the line network is pairwise Nash stable. It is easy to see that its efficiency ratio is linear in n, because its social cost is $O(n^3)$ while the social cost of an efficient network is $O(n^2)$. Thus the *price of anarchy* (as defined by Papadimitriou in [26]) is at least linear in n. Another consequence is that all efficient networks are pairwise Nash stable, whenever $\pi > c_{\min}$; thus it is important to try to select good equilibria (in terms of efficiency). The dynamics we consider are well matched to this purpose.

Let $\ell > 1$ be a given integer. We study discrete-time dynamics that proceeds in *rounds*. At each round k, an exogenous process (called an *activation process*) selects an active node $u_k \in V$. An important case that we study is the *uniform activation process*: at each round k, the active node is selected independently and uniformly at random from the set of nodes. Thus, under the uniform activation process, for all $k > 0$ and all $v \in V$, $\mathbb{P}[u_k = v] = 1/n$ independent of the past history of the activation process.

Let $\left(G^{(k)}, \mathbf{P}^{(k)}, \Gamma^{(k)}\right)$ be the state of the network at the beginning of round k. The dynamics at round k proceeds in two *stages*.

During the first stage, the active node u_k selects a set of contracts (possibly empty) from $\Gamma^{(k)}$ it currently participates in, and removes them. All payments associated to those contracts are set to zero. If all contracts associated to an edge $u_k x \in G^{(k)}$ are removed, then the edge $u_k x$ is removed. Let $\left(G_1^{(k)}, \Gamma_1^{(k)}, \mathbf{P}_1^{(k)}\right)$ be the resulting state of the network following stage 1 of round k.

During the second stage, the active node u_k either selects a new set of contracts (possibly empty) from $\Gamma_1^{(k)}$ it participates in and remove them, or it selects a node w in its ℓ-*neighborhood* from $G^{(k)}$, i.e., from among those nodes such that $d(u_k, w; G^{(k)}) \leq \ell$). In this case u_k proposes the contract (u_k, w) to w. If w accepts, the contract (u_k, w) is added to $\Gamma_1^{(k)}$, the edge $u_k w$ is added to $G_1^{(k)}$, and we set

$$p_{u_k w} = Q(u_k, w; G_1^{(k)} + u_k w).$$

Note that the active node only contemplates deviating with nodes in its ℓ-neighborhood, thus making our equilibrium selection process both *decentralized* and *local*.

We assume the dynamics are *myopic* in the following sense:

- the active node u_k selects its actions in order to maximize its utility at the end of the round; and
- the node w selected during the second stage accepts or rejects the contract in order to maximize its utility given the state at the end of the first stage, i.e., given $\left(G_1^{(k)}, \Gamma_1^{(k)}, \mathbf{P}_1^{(k)}\right)$.

As a tie-breaking rule, we assume the following notion of "inertia."

Assumption 3 (Inertia). *Let u_k be the active node at round k, and let $\left(G_1^{(k)}, \mathbf{P}_1^{(k)}, \Gamma_1^{(k)}\right)$ be the state of the network after the first stage of round k. Let $W \subseteq V$ be the subset of nodes to whom u_k considers offering a contract at the second stage, i.e., such that the utility of u_k is maximized after the second stage. If $|W| > 1$, then u_k selects the node in W it was most recently connected to. If no such node exists, u_k picks one uniformly at random.*

Assumption 3 states that, if the active node has more than one optimal choice after the first stage, it will choose to deviate with the node it was most recently connected to.

A tie-breaking rule is necessary as the active node at a round, say u_k, may not have a unique utility-maximizing choice of a "partner" node at stage 2. Thus, in order to avoid oscillations induced by the possibility of multiple optimal choices, a tie-breaking rule must be assumed. While we have chosen a specific notion of inertia, we emphasize that many other assumptions can also lead to convergent dynamics. For instance, among utility-maximizing choices of w, if node u_k always chooses the node w with the highest degree, our convergence results remain valid.

Convergence of our dynamics is defined as follows.

Definition 3 (Convergence). *Given an initial state $\left(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)}\right)$, we say that the dynamics converge if, almost surely, there exists $K > 0$ such that, for all $k > K$,*

$$\left(G^{(k)}, \mathbf{P}^{(k)}, \Gamma^{(k)}\right) = \left(G^{(k+1)}, \mathbf{P}^{(k+1)}, \Gamma^{(k+1)}\right).$$

7 Results

In this section we state and prove our main results. We interpret these results by analyzing the efficiency of the limiting topologies in our dynamics.

Theorem 2 (Convergence Theorem for $\pi > 0$). *Let $\ell \geq 2$ be given. Suppose that Assumptions 1 to 3 hold. Further, assume that $\pi > 0$.*

Let $\left(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)}\right)$ be an outcome of the static game such that $G^{(0)}$ is connected. Assume that the activation process is such that, almost surely, all nodes are activated infinitely often. Then the dynamics started at $\left(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)}\right)$ converge. Further, if the activation process is uniform, the convergence time is polynomial.

For a given realization of the activation process, let $(G, \mathbf{P}, \Gamma)$ be the limiting state. Then:

1. *G is a tree where all internal (i.e., non-leaf) nodes are of minimum routing cost; and*
2. *the limiting state $(G, \mathbf{P}, \Gamma)$ is pairwise Nash stable.*

In order to state our result when $\pi = 0$, we need two extra tie-breaking assumptions. The idea is that $\pi = 0$ can induce bilateral deviations where the increase in cost for adding a new edge is nil. For such deviations, we want the value of the corresponding contracts to be zero. This is consistent with our interpretation of contracts as a way incentivize nodes to accept connections even when their cost in the network increases. Thus, if the cost of adding a connection is zero, there should be no need for an incentive. That is the content of the following assumption.

Assumption 4 (Zero Value Contracts). *We assume that the contracting function yields a zero utility transfer if and only if there is no extra cost associated to adding the proposed contract. More formally, for all distinct nodes u and v, for all network topologies G, $Q(u, v; G) = 0$ if and only if $C(v; G + uv) = C(v, G)$.*

We now need to decide whether the active node should contemplate adding zero-value contracts. This is a tie-breaking rule as contracts can only be added during the second stage of the dynamics, and thus the utility of the active node would stay constant should it propose a zero-value contract. As a tie breaking rule, we assume that the active node *would not* propose a zero value contract.[2]

Assumption 5 (Dynamics of Zero Value Contracts). *Let u_k be the active node, and let $(G, \mathbf{P}, \Gamma)$ be the state of the network prior to the bilateral deviation considered. If the utility of u_k after successfully adding the link $u_k w$ (for any such w) is identical to that when in state $(G, \mathbf{P}, \Gamma)$, then u_k does not propose any contract to w.*

The motivation for Assumption 5 is that even if a contract has zero value, in reality there is some implicit "burden" to setting up a contract, so that establishing a contract

[2] Note that if a node is *proposed* a zero value contract we assume that it would accept it. The rationale behind Assumption 5 is that the active node can choose what node to bilaterally deviate with, which is not the case for a node being proposed a zero-value contract.

with no value is not desirable. Assumption 5, together with Assumption 4, implies that if the best action for the active node is to add a zero value contract, it prefers to pass.

In order to prove our second theorem, we also need to assume that $\ell \geq 3$, i.e. we allow nodes to be aware of their 3-neighborhood when considering their second stage deviation (recall that Theorem 2 requires only $\ell \geq 2$).

Theorem 3 (Convergence Theorem for $\pi = 0$). *Let $\ell \geq 3$ be given. Assume that Assumptions 1 to 5 hold. Further, assume that $\pi = 0$.*

Let $\left(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)}\right)$ be an outcome of the static game such that $G^{(0)}$ is connected. Assume that the activation process is such that, almost surely, all nodes are activated infinitely often. Then the dynamics started at $\left(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)}\right)$ converge. Further, if the activation process is uniform, the convergence time is polynomial.

For a given realization of the activation process, let $(G, \mathbf{P}, \Gamma)$ be the limiting state. Then the following hold:

1. *Let G' be the graph obtained by contracting all cliques in G and replacing them with a node whose per-unit routing cost is set to be that of the smallest routing cost in that clique. Then,*
   (a) G' is a tree;
   (b) all internal nodes have minimum per-unit routing cost.
2. *For any clique in G, only nodes with the smallest per-unit routing cost (among nodes of the same clique) can have edges to nodes outside of the clique.*
3. *The limiting state $(G, \mathbf{P}, \Gamma)$ is pairwise Nash stable.*

Upon inspection of the dynamics considered in this paper, a generalization one might consider is to allow the active node to unilaterally or bilaterally deviate during *both* stages (recall that our dynamics allow only unilateral deviations during the first stage). However, as shown in [21], such dynamics may fail to converge under our assumptions.

8 Discussion of Results

In this section we provide a brief analysis of the results of Theorem 2, and a formal analysis of the results of Theorem 3.

In our previous work [4], under the same assumptions, we proved that a significantly restricted version of the dynamics considered in this paper converge to the same set of outcomes. In particular, in this paper activated nodes choose which links to break in the first stage of each round, whereas our model in [4] imposed random breaking of edges. Further, this paper analyzes a setting where nodes may be limited in their visibility of the network, and thus only contract with nodes in their ℓ-neighborhood; by contrast, the dynamics in [4] requires nodes be able to "see" the entire network (i.e., that $\ell = n$). In addition, we provide a novel proof technique via a Lyapunov argument that succinctly addresses these generalizations.

As argued in [4], our dynamics select good equilibria in that only nodes of minimum routing cost forward packets. An important special case is where there is a unique node i of minimum per-unit routing cost, i.e., such that $c_i < c_j$ for all $j \neq i$. In that case, our dynamics converge to a star centered around node i, which is the most efficient pairwise stable outcome. In [27, 10], the term *price of stability* was coined to refer to the ratio of the efficiency of the best equilibrium to the optimum efficiency; thus our dynamics select an equilibrium that achieves the price of stability.

We now consider the case where $\pi = 0$ (cf. Theorem 3), with a view towards demonstrating that significantly different results are obtained when compared with the case where $\pi > 0$. First, we note that when $\pi = 0$, any outcome of the static game such that the network topology is the complete network is pairwise Nash stable. To see this, it suffices to note that when the network topology is the complete network, no node forwards any traffic, and thus the cost to all nodes is minimized. Next, by Assumption 4, it follows that all contracts in such network topology have a zero payment associated with them. Thus no unilateral deviation is profitable, and no bilateral deviation is possible as all edges are already part of the network. Recall that in [25] it is proved that, when $\pi < c_{\min}$, the only efficient outcomes of the static game are those where the network topology is the complete network. Thus, when $\pi = 0$, we conclude that *any efficient outcome is pairwise Nash stable.*

Surprisingly, allowing $\pi = 0$ can lead both to situations where *more* efficient outcomes are chosen than when $\pi > 0$, as well as situations where *less* efficient outcomes are chosen than when $\pi > 0$. First, from Theorem 2, note that when $0 < \pi < c_{\min}$, our dynamics *cannot* select an efficient outcome. By contrast, in [21] we select a set of parameters of the model, and construct a contracting function such that the efficient outcome is a fixed point of our dynamics when $\pi = 0$. Thus, by allowing $\pi = 0$, we can select more efficient outcomes than when $\pi > 0$.

However, allowing $\pi = 0$ can also make the *most inefficient* outcome selected worse. This is also shown in [21]. where we select a set of parameters, and construct a contracting function such that the social cost of a fixed point when $\pi = 0$ is strictly larger than the social cost of *any* fixed point using the same parameters and contracting function when $0 < \pi < c_{\min}$. Thus, in this setting, by allowing $\pi = 0$ we can select *less* efficient outcomes than when $\pi > 0$.

Acknowledgments

The authors benefited from helpful conversations with John N. Tsitsiklis and Matthew O. Jackson. This work was supported by DARPA under the ITMANET Program, by the Canada Research Chairs Program, by the NSERC, and by the NSF under grant CMMI-0620811.

References

[1] Jackson, M.O.: A survey of models of network formation: Stability and efficiency. Working Paper 1161, California Institute of Technology, Division of the Humanities and Social Sciences (2003)
[2] Jackson, M.O.: The stability and efficiency of economic and social networks. Microeconomics 0211011, EconWPA (November, 2002)
[3] Jackson, M.O., Wolinsky, A.: A strategic model of social and economic networks. Journal of Economic Theory 71(1), 44–74 (1996)

[4] Arcaute, E., Johari, R., Mannor, S.: Network formation: Bilateral contracting and myopic dynamics. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 191–207. Springer, Heidelberg (2007)
[5] Johari, R., Mannor, S., Tsitsiklis, J.N.: A contract-based model for directed network formation. Games and Economic Behavior 56(2), 201–224 (2006)
[6] Jackson, M.O., Watts, A.: The existence of pairwise stable networks. Seoul Journal of Economics 14(3), 299–321 (2001)
[7] Jackson, M.O., Watts, A.: The evolution of social and economic networks. Journal of Economic Theory 106(2), 265–295 (2002)
[8] Roughgarden, T.: Selfish Routing and the Price of Anarchy. MIT Press, Cambridge (2005)
[9] Lin Chen, H., Roughgarden, T., Valiant, G.: Designing networks with good equilibria. In: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007) (2008)
[10] Anshelevich, E., Dasgupta, A., Tardos, É., Wexler, T.: Near-optimal network design with selfish agents. In: Proceedings of ACM Symposium on the Theory of Computing, pp. 511–520 (2003)
[11] Fiat, A., Kaplan, H., Levy, M., Olonetsky, S., Shabo, R.: On the price of stability for designing undirected networks with fair cost allocations. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 608–618. Springer, Heidelberg (2006)
[12] Fudenberg, D., Levine, D.K.: The Theory of Learning in Games. MIT Press, Cambridge (1998)
[13] Cesa-Bianchi, N., Lugosi, G.: Prediction, Learning, and Games. Cambridge University Press, Cambridge (2004)
[14] Mannor, S., Shamma, J.: Multi-agent learning for engineers. Artificial Intelligence, 417–422 (2007); Special Issue on Foundations of Multi-Agent Learning
[15] Altman, E., Boulogne, T., El-Azouzi, R., Jiménez, T., Wynter, L.: A survey on networking games in telecommunications. Computers & Operations Research 33(2), 286–311 (2006)
[16] Falkner, M., Devetsikiotis, M., Lambadaris, I.: An overview of pricing concepts for broadband IP networks. IEEE Communications Surveys 3(2) (2000)
[17] Briscoe, B., Darlagiannis, V., Heckman, O., Oliver, H., Siris, V., Songhurst, D., Stiller, B.: A market managed multiservice Internet (M3I). Computer Communications 26(4), 404–414 (2003)
[18] Komali, R., MacKenzie, A.: Distributed topology control in ad hoc networks: A game theoretic perspective. In: Proceedings of IEEE CCNC, pp. 563–568 (2006)
[19] Eidenbenz, S., Kumar, V., Zust, S.: Equilibria in topology control games for ad hoc networks. In: Proceedings of International Conference on Mobile Computing and Networking, pp. 2–11 (2003)
[20] Fabrikant, A., Luthra, A., Maneva, E.N., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: Proceedings of the Twenty-Second ACM Symposium on Principles of Distributed Computing (PODC 2003), Boston, Massachusetts, USA, July 13-16, 2003, pp. 347–351 (2003)
[21] Arcaute, E., Johari, R., Mannor, S.: Local two-stage myopic dynamics for network formation games. Technical report, Stanford University Management Science and Engineering (2008)
[22] Osborne, M.J., Rubinstein, A.: Bargaining and Markets. Academic Press, San Diego (1990)
[23] Bloch, F., Jackson, M.O.: The formation of networks with transfers among players. Journal of Economic Theory 133(1), 83–110 (2007)

[24] Corbo, J., Parkes, D.: The price of selfish behavior in bilateral network formation. In: Proceedings of the Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing, PODC 2005, Las Vegas, NV, USA, July 17-20, 2005, pp. 99–107 (2005)

[25] Arcaute, E., Johari, R., Mannor, S.: Network formation: Bilateral contracting and myopic dynamics. Technical report, Stanford University Management Science and Engineering (2007)

[26] Papadimitriou, C.H.: Algorithms, games, and the internet. In: Proceedings on 33rd Annual ACM Symposium on Theory of Computing, Heraklion, Crete, Greece, July 6-8, 2001, pp. 749–753 (2001)

[27] Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: 45th Symposium on Foundations of Computer Science (FOCS 2004), Rome, Italy, 17-19 October 2004, pp. 295–304 (2004)