Network Formation:
Bilateral Contracting and Myopic Dynamics

Esteban Arcaute (arcaute@stanford.edu)
Ramesh Johari (rjohari@stanford.edu)
Shie Mannor (shie.mannor@mcgill.ca)

July 25, 2007

Abstract

We consider a network formation game where nodes wish to send traffic to each other. Nodes contract bilaterally with each other to form bidirectional communication links; once the network is formed, traffic is routed along shortest paths (if possible). Cost is incurred to a node from four sources: (1) routing traffic; (2) maintaining links to other nodes; (3) disconnection from destinations the node wishes to reach; and (4) payments made to other nodes. We assume that a network is stable if no single node wishes to unilaterally deviate, and no pair of nodes can profitably deviate together (a variation on the notion of pairwise stability). We study such a game under a form of *myopic best response dynamics*. In choosing their action, nodes optimize their single period payoff only. We characterize a simple set of assumptions under which these dynamics converge to a stable network; we also characterize an important special case, where the dynamics converge to a star centered at a node with minimum cost for routing traffic. In this sense, our dynamics naturally *select* an efficient equilibrium.

1 Introduction

Modern communication networks operate at a scale that makes the use of centralized policies/protocols challenging. It is therefore natural to study networks that are formed by ad-hoc and decentralized dynamics. The influence of the behavior of individual nodes on the global operation of the network is among the main challenges in analyzing such large systems. The traditional approach in systems and game theory is to consider each node as a strategic agent whose utility function dictates his behavior. This typically leads to describing the strategic behavior of nodes as an optimization problem and considering some equilibrium concept (Nash, correlated, Wardrop, etc.) as the "desired" outcome of the system. From the perspective of system design, significant research has been devoted to analyzing properties of the equilibrium concept that are important to the network designer, particularly social welfare. In many instances the set of equilibria can be quite large, so even converging to an equilibrium is in itself non-trivial; this only hints at the difficulty of designing dynamics that select a "good" equilibrium. As shown in [8, 2, 22], there is no reason to expect equilibria to be efficient. Indeed, in [17] a network formation game is studied where inefficiency grows linearly in the size of the network.

The focus of this paper is on the dynamics of networks of multiple self-optimizing agents. Instead of focusing on one of the standard static equilibrium concepts we focus on natural myopic dynamics; motivation for our approach can be found in Arrow's statement that "the attainment of equilibrium requires a disequilibrium process" [4]. In this paper we consider such a disequilibrium process for a class of games that naturally arise from ad-hoc networks. While there are dynamic notions of equilibrium that have been developed for multi-stage dynamic game models (such as subgame perfect and sequential equilibria), we do not consider those here for two reasons. First, the information required from each player is typically quite significant in such equilibrium concepts, as is the knowledge each player must have about other players' intentions. Second, the complexity of finding an optimal policy in such models is often unrealistic for ad-hoc networks.

Our focus is on *network formation games* (NFGs). These games describe the interaction between a collection of nodes that wish to form a network. Such models have been introduced and studied in the economics literature; see, e.g., [5, 16, 13]. Each of the nodes in the network is a decision maker and a network is formed through interaction between the decision makers. We are interested in understanding and characterizing the networks that result when the decision makers interact to choose their connections. In particular, we will focus on the role of bilateral contracting and the dynamic process of network formation in shaping the network structure. Following standard game theory terminology, we refer to the decision makers as *players*.

Network formation games capture the following properties:

1. Nodes in the network are strategic agents;

2. Links in the network represent bilateral agreements between their endpoints;

3. Nodes choose whom to accept connections from and whom to connect to; and

4. Given a network topology, nodes' payoffs are obtained from two contributions: (1) intrinsic utility extracted from participating in the network; and (2) transfers of utility, or "payments", between nodes participating in the same link.

For static NFGs, we consider a natural solution concept called *pairwise stability*. Roughly speaking, a network is pairwise stable if no coalition of at most 2 nodes can profitably deviate. We eschew the more common approach of studying Nash equilibria, as Nash equilibria are typically either trivial or fail to exist (depending on the model formulation); see, e.g., [16]. In order to account for our interpretation of links as contractual agreements, we assume that the value of all contracts stays constant in any deviation unless those contracts were explicitly modified during the deviation.

From a system or network design perspective there are several benefits to consider the NFG framework. First, viewing nodes as strategic agents facilitates reasoning in terms of optimizing a utility function. Second, the utility functions are driven by both the global topology of the network and by local interactions. Lastly, nodes in networks are typically only aware of local interactions and can only be connected through local connections. Thus, pairwise stability which focuses on a single connection is easier to analyze than other concepts of stability that consider several connections simultaneously.

It is natural to ask how a network would *evolve* under some dynamics in light of the underlying network formation game. One of the main contributions of this paper is to define a two-stage dynamics that naturally fits the solution concept of pairwise stability. These dynamics describe the myopic behavior of a strategic agent, should it be allowed to deviate within a particular set of allowable deviations. Specifically, given a current network configuration, a node u contemplating a pairwise deviation with, say, node v would *first* unilaterally deviate from its current action, and *later* engage in a bilateral deviation with node v; these two stages comprise one time step in the dynamics we define. The intuition behind such behavior is that if node u decides to engage directly in a bilateral deviation with node v, then node u might not be able to perform all the actions he wanted that did not require node v's approval. Note that complex deviations involving coalitions of multiple agents are not allowed, consistent with the equilibrium notion of pairwise stability.

The second main contribution of this paper is the characterization of the payment structures (called *contracting functions*) that lead to convergent dynamics. Loosely speaking, the contracting functions should have a *monotonicity* property, which informally assumes that the utility transfer should increase as the burden to the node considered increases; and an *anti-symmetry* property, which requires that the value of a contract is independent of which of the two agents initiated the contract. We also show that if all links have an intrinsic (random) "expiration date," i.e., if all links are broken at random infinitely often, then only the monotonicity property is required to ensure convergence.

Our third key contribution is to establish that convergence is assured to a "good" equilibrium for the dynamics we consider. We prove that, in a special case, the proposed dynamics converge to the most efficient pairwise stable network. Furthermore, under additional conditions the expected convergence time is polynomial in the number of nodes.

Our results are established in the context of a natural game theoretic model inspired by traffic routing in networks. In the model we consider, nodes incur a cost due to three sources: (1) routing traffic; (2) maintaining links to other nodes; and (3) disconnection from destinations the node wishes to reach. Such a model is an appropriate reference point for ad-hoc network creation in data communications contexts.

Our work touches on several related threads of the literature. Most closely related is the work on network formation games in economics (see [13] for a survey). In particular, the work of Jackson and Watts also considers dynamics for network formation games [14, 15], for a utility model that is unrelated to ours; while in their dynamics any unilateral or bilateral deviation may occur in a single stage, our dynamics are designed so that each stage consists of a unilateral deviation *followed* by a bilateral deviation. It is this latter property that allows us to *select* desirable equilibria. Note that, when all links have (random) "expiration dates", the assumption that links are broken at random infinitely often is similar to the stochastic dynamics considered in [15], where the decision on the link considered during the round is reversed exogenously with some probability. Despite the dissimilarities in cost structure, the use of randomness improves convergence results both in our setting and in that of [15].

Our work is also related to the literature on learning in games. (See, e.g., [12, 7, 19] for surveys.) This literature has recently benefited from the application of control techniques; see, e.g., [23]. In this approach, typically 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. Motivated by design of networks, our approach departs significantly from this literature, as we are also 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., [1] for a survey, and [11, 6] 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, 9, 10]. However, these works all consider Nash equilibria, whereas our focus is on *pairwise* interactions between nodes.

The remainder of the paper is organized as follows. In Section 2, we describe the game model we consider, including routing costs, link maintenance costs, disconnectivity costs, and monetary transfers between nodes. In Section 3 we precisely define the notion of stability and efficiency. In Section 4 we specialize the model to the traffic routing game of the preceding paragraph. In Section 5 we present the proposed dynamics and the related convergence results. Some proofs are deferred to online appendices in the techical report [3].

2 The Formation Game

In this section, we formally present the NFG we consider. The players are the set of nodes of the network. Nodes receive a reward that depends on the network topology that arises from the formation. The game models a scenario where each link in the network is the result of a bilateral "contract" between its end nodes. Each contract has a node seeking the agreement and a node accepting it, and therefore there is some utility transfer from the seeking node to the accepting one. In addition to the utility transfers, nodes obtain utility that depends on the topology created. In the rest of the section we make this explanation formal by prescribing the actions available to the players and their utility function.

We use the notation $G = (V, E)$ to denote a graph, or *network topology*, consisting of a set of n nodes V and edges E; the nodes will be the players in the NFG. We assume throughout that all edges in G are *undirected*; we use ij to denote an undirected edge between i and j. We will typically use the shorthand $ij \in G$ when the edge ij is present in E. We use $G + ij$ and $G - ij$ to denote, respectively, adding and subtracting the link ij to the graph G.

For a node $i \in V$, let $v_i(G)$ be the monetary value to node i of network topology G. Let P_{ij} denote a payment from i to j; we assume that if no undirected link ij exists, or if $i = j$, then $P_{ij} = 0$. We refer to $\mathbf{P} = (P_{ij}, i, j \in V)$ as the *payment matrix*. Given a payment matrix P, the total transfer of utility to node i is the sum of payments received by i minus the sum of payments made by i, that is: $T_i(\mathbf{P}) = \sum_j P_{ji} - P_{ij}$. Thus the total utility of node i in graph G is:

$$U_i(\mathbf{P}, G) = T_i(\mathbf{P}) + v_i(G).$$

We consider a network formation game where each node selects other nodes they wish to connect to, as well as those they are 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 $ij \in G(\mathbf{T}, \mathbf{F})$.

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

The contracting graph Γ and the network topology G determine the transfers between the nodes. Formally, we assume that there is a *contracting function* $Q(i, j; G)$ that gives the payment in a contract from i

to j when the network topology is G; note that if $Q(i,j;G)$ is negative, then j pays i. Given strategy vectors $\mathbf{T}$ and $\mathbf{F}$, the payment matrix $\mathbf{P}(\mathbf{T},\mathbf{F})$ at the outcome of the game is given by:

$$P_{ij}(\mathbf{T},\mathbf{F}) = \begin{cases} Q(i,j;G(\mathbf{T},\mathbf{F})), & \text{if } (i,j) \in \Gamma(\mathbf{T},\mathbf{F}); \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

Thus given strategy vectors $\mathbf{T}$ and $\mathbf{F}$, the payoff to node i is $U_i(G(\mathbf{T},\mathbf{F}),\mathbf{P}(\mathbf{T},\mathbf{F}))$. By an abuse of notation, 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 $(G,\Gamma,\mathbf{P})$.

There are two plausible interpretations of the contracting function. First, a system designer or an external regulator could determine that nodes must have pre-negotiated tariffs which are encoded in the contracting function. In this case, the value of the contract changes 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 a negotiation process. 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 currently 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 [21]. We investigate this example in further detail in Appendix A of [3].

The contracting functions allow considerable design flexibility in a distributed setting. Instead of focusing on a particular contracting function, we will be interested in contracting functions exhibiting two natural properties: *monotonicity* and *anti-symmetry*. We start with some additional notation: given $j \neq i$, define the *difference* in cost to node i between graph G and graph $G+ij$ as $\Delta C_i(G,ij) = C_i(G+ij) - C_i(G)$. (Note that if $ij \in G$, then $\Delta C_i(G,ij) = 0$.)

Property 1 (Monotonicity) *Let G be a graph such that $ij \notin G$ and $ik \notin G$. We say that the contracting function is* monotone *if:*

$$\Delta C_j(G,ij) > \Delta C_k(G,ik) \textit{ if and only if } Q(i,j;G+ij) > Q(i,k;G+ik).$$

(Note that since j and k are interchangeable, if the differences on the left hand side of the previous definition are equal, then the contract values on the right hand side must be equal as well.) Informally, monotonicity requires that the payment to form a link must increase as the burden of forming that link increases on the accepting node.

The second property is motivated by the observation that, in general, $Q(i, j; G)$ is not related to $Q(j, i; G)$; anti-symmetry asserts these values must be equal.

Property 2 (Anti-symmetry) *We say that the contracting function Q is* anti-symmetric *if, for all nodes i and j, and for all graphs G, we have $Q(i, j; G) = -Q(j, i; G)$.*

A contracting function that is anti-symmetric has the property that at any feasible outcome of the game, the payment for a link ij does not depend on which node asked for the connection.

3 Stability and Efficiency

We study our game through two complementary notions: stability and efficiency. Since nodes act as self-interested players, we define a game-theoretic notion of equilibrium, called *pairwise stability* (first introduced by Jackson and Wolinsky [16]). Informally, pairwise stability requires that no unilateral deviations by a single node are profitable, and that no bilateral deviations by any pair of nodes are profitable. Since we are also interested in system-wide performance from a global perspective we study the *efficiency* of the network as well. We measure the efficiency of a network topology via the total value obtained by all nodes using that topology.

We start by considering game theoretic notions of equilibrium for our model. The simplest notion of equilibrium is the celebrated *Nash equilibrium*: a strategy profile $(\mathbf{T}, \mathbf{F})$ is a Nash equilibrium if no node i can execute a profitable *unilateral deviation*, i.e., strictly improve its payoff by altering either or both of the sets T_i and F_i. However, Nash equilibrium lacks sufficient predictive power in our model due to the presence of trivial equilibria. For example, it is not hard to see that $F_i = T_i = \emptyset$ is a Nash equilibrium regardless of the cost structure or contracting function: no node can affect the outcome through a unilateral deviation, so no unilateral deviation is profitable. The inadequacy of the Nash equilibrium as a solution concept is well known for NFGs; see [13].

Link formation is inherently *bilateral*: the consent of two nodes is required to form a single link. For this reason we consider a notion of stability that is robust to *bilateral* deviations, known as *pairwise stability*. Informally, pairwise stability of a strategy vector requires that both no unilateral deviations are profitable and that no two nodes can collude to improve their payoff.

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. 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. Any deviation will of course change both the network topology, as well as the contract graph. 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. The formal definition of pairwise stability follows; note that it is similar in spirit to the definition of Jackson and Wolinsky [16].

Definition 1 *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:*

$$P'_{k\ell} = \begin{cases} P_{k\ell}, & \textit{if } (k,\ell) \in \Gamma' \textit{ and } (k,\ell) \in \Gamma; \\ Q(k,\ell;G'), & \textit{if } (k,\ell) \in \Gamma' \textit{ and } (k,\ell) \notin \Gamma; \\ 0, & \textit{otherwise}. \end{cases} \tag{2}$$

Then $(\mathbf{T}, \mathbf{F})$ is a pairwise stable equilibrium *if: (1) No unilateral deviation is profitable, i.e., for all i, and for all $\mathbf{T}'$ and $\mathbf{F}'$ 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 $\mathbf{T}$ and $\mathbf{F}$ 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 (2) is a formalization of the discussion above. When nodes i and j deviate to the strategy vectors $\mathbf{T}'$ and $\mathbf{F}'$, all payments associated to preexisting contracts remain the same. If a contract is formed, the payment becomes the value of the contracting function given the new graph. Finally, if a contract is broken, the payment of course becomes zero. These conditions give rise to the new payment matrix $\mathbf{P}'$. Nodes then evaluate their payoffs before and after a deviation. 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 be interested in pairwise stability of the network topology and contracting graph, rather than pairwise stability of strategy vectors. We will thus say that a feasible outcome $(G, \Gamma, \mathbf{P})$ is a *pairwise 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 stable equilibrium; and (2) $(\mathbf{T}, \mathbf{F})$ give rise to $(G, \Gamma, \mathbf{P})$. Note that for all i and j such that $ij \in G$ we must have $P_{ij} = Q(i, j; G)$ in a pairwise stable outcome.

The following lemma yields a useful property of pairwise stable outcomes; the proof can be found in [3].

Lemma 1 *Let $(G, \Gamma, \mathbf{P})$ be a pairwise stable outcome. Then for all nodes i and j, if $(i, j) \in \Gamma$ and $(j, i) \in \Gamma$, then $Q(i, j; G) = 0$ and $Q(j, i; G) = 0$.*

We will investigate the *efficiency* of pairwise stable equilibria. 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 [20], it is not hard to show that a feasible outcome $(G, \Gamma, \mathbf{P})$ is Pareto efficient if and only if $G \in \arg\min_{G'} S(G')$, where $S(G)$ is the *social cost function*:

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

(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.) Given a graph G, we define the *efficiency* of G as the ratio $S(G)/S(G_{\text{eff}})$, where G_{eff} is the network topology in a Pareto efficient outcome.

4 A Traffic Routing Utility Model

Motivated by highly reconfigurable wireless ad hoc communication networks, we specify an NFG model where the nodes' utility depends on the traffic that is routed through the network. In particular, nodes extract utility per unit of data they successfully send through the network and experience per-unit routing costs when in the data network, as well as maintenance costs per adjacent link.

We first describe the traffic routing model. Formally, we assume that each user i wants to send one unit of traffic to each node in the network; we refer to this as a *uniform all-to-all* traffic matrix. We assume that given a network topology, traffic is routed along shortest paths, where the length of a path is measured by the number of hops. In case of multiple shortest paths of equal length, traffic is split equally among all available paths.

We assume that each node experiences three types of costs:

1. *Routing costs.* Let $f_i(G)$ be the total traffic that transits through i plus the total traffic received by i. We assume that node i experiences a positive routing cost of c_i per unit of traffic. Thus given a graph G, the total routing cost experienced by node i is $R_i(G) = c_i f_i(G)$.

2. *Link maintenance cost.* A maintenance cost of $\pi > 0$ is incurred by the endpoints of each link (the effective cost of a single link is 2π). [1] Thus given a graph $G = (V, E)$, the total link maintenance cost incurred by node i is $M_i(G) = \pi d_i(G)$, where $d_i(G)$ is the degree of node i in the graph G.

3. *Disconnection cost.* We assume that each node experiences a cost of $\lambda > 0$ per unit of traffic not sent because the network is not connected. [2] Thus given a graph G, the cost to a node i from incomplete connectivity, or disconnection cost, is $\lambda(n - n_i(G))$, where $n_i(G)$ is the number of nodes i can reach in the graph G.

Thus the total cost to a node i in a graph G is:

$$C_i(G) = R_i(G) + M_i(G) + D_i(G). \tag{3}$$

4.1 Pairwise Stability

We now characterize pairwise stable outcomes, given the cost model (3). We start with the following structural characterization; the proof can be found in [3].

Proposition 1 *Let $(G, \Gamma, \mathbf{P})$ be a pairwise stable outcome. Then G is a forest (i.e., all connected components of G are trees).*

The preceding proposition shows the "minimality" of pairwise stable graphs: since our payoff model does not include any value for redundant links, any pairwise stable equilibria must be forests. An interesting open direction for our model includes the addition of a utility for redundancy (e.g., for robustness to failures).

Most of the pairwise stable equilibria we discuss are framed under the following assumption on the disconnectivity cost λ.

Assumption 1 (Disconnection Cost) *Given a contracting function Q, the disconnectivity cost $\lambda > 0$ is such that for all disconnected graphs G and for all pairs i and j that are disconnected in G, there holds $\Delta C_i(G, ij) + Q(i, j; G + ij) < 0$ and $\Delta C_i(G, ij) - Q(j, i; G + ij) < 0$.*

This implies that if nodes i and j are not connected in G, then both are better off by forming the link ij using either the contract (i, j) or (j, i). (Note that if Q is anti-symmetric the second condition is trivially satisfied.)

[1] The link maintenance cost does not depend on the identities of the endpoints of the link; this homogeneity assumption is made for technical simplicity.

[2] The parameter λ is identical for all nodes; again, this homogeneity assumption simplifies the technical development.

The preceding assumption is meant to ensure that we can restrict attention to connected graphs in our analysis. From our utility structure, it is easy to see that only the payments and disconnectivity costs act as incentives to nodes to build a connected network topology. But payments alone are not enough to induce connectivity, since of course the node paying for a link feels a negative incentive due to the payment. We emphasize that the preceding assumption is made assuming that *the contracting function and all other model parameters are given*, so that the threshold value of λ necessary to satisfy the preceding assumption may depend on these other parameters. Nevertheless, as we will see this assumption has interesting implications for our model. It is clear from our model that if all other model parameters are fixed, then a λ satisfying the preceding assumption must exist. Examples where λ scales as $O(n)$ can be found in Appendix B of [3].

If Assumption 1 holds, we have the following corollary about pairwise stable outcomes; the proof is immediate.

Corollary 1 *If Assumption 1 holds, all pairwise stable outcomes are trees.*

From the preceding corollary, we can prove the following simple characterization of pairwise stable outcomes; see [3] for the proof.

Proposition 2 *Suppose that Assumption 1 holds, and that Q is monotone. Let $(G, \Gamma, \mathbf{P})$ be a feasible outcome where G is a tree. Then $(G, \Gamma, \mathbf{P})$ is pairwise stable if and only if no pair of nodes can profitably deviate by simultaneously breaking one link and forming another, i.e.: given nodes i and j and any link $ik \in G$, let $G = G - ik + ij$, $\Gamma' = (\Gamma \setminus \{(i,k),(k,i)\}) \bigcup \{(i,j)\}$, and define $\mathbf{P}'$ as in* (2)*. Then:*

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

4.2 Efficiency of Equilibria

Pairwise stable equilibria are in general inefficient, as the following example suggests.

Example 1 Suppose Assumption 1 holds, and assume that the contracting function is monotone. We assume there is a unique node $u_{\min}$ such that $c_{u_{\min}} < c_i$ for all nodes $i \neq u_{\min}$; let $c_{\min} = c_{u_{\min}} > 0$. For simplicity, assume that all other nodes i have the same per unit routing cost $c_i = c > c_{min}$. Let S be a star centered at $w \neq u_{min}$, and let the contracting graph be Γ_S such that for all $v \neq w$, we have $(w, v) \in \Gamma_S$ (and these are the only contracts in Γ_S). Let $\mathbf{P}$ be the resulting payment matrix. Suppose Q is such that for all $v \neq w, u_{\min}$, $Q(v, u_{min}; S - vw + vu_{min}) < 0$. Since S is a tree, we can use Proposition 2 to prove that $(S, \Gamma_S, \mathbf{P})$ is pairwise stable; see [3] for details.

It is clear that the preceding graph cannot be efficient; indeed, it is not even efficient among all trees: a star centered at $u_{\min}$ would generate lower social cost than a star centered at any other network topology. As long as the contracting function is monotone, it is possible to show that any tree where non-leaf nodes have minimum routing cost can be sustained as pairwise stable equilibrium. This is the content of the next proposition.

Proposition 3 *Suppose that Assumption 1 holds. Let $(G, \Gamma, \mathbf{P})$ be a feasible outcome such that G is a tree, and any non-leaf node i has $c_i = \min_j c_j$; i.e., all internal nodes of G have minimum per-unit routing cost. Then $(G, \Gamma, \mathbf{P})$ is pairwise stable.*

Proof. From Proposition 2, it is sufficient to check whether given $(G, \Gamma, \mathbf{P})$, it is not profitable for any pair of nodes to delete an edge and add another edge. Assume that u, v and w are such that $uv \in G, uw \notin G$. There are two possibilities: if v is not on the path from u to w in G, then removing the edge uv from G disconnects the graph. Given Assumption 1, this cannot be profitable. On the other hand, suppose the path from u to w passes through v. This implies that $G - uv + uw$ is a tree, and that v is an internal node; hence v has minimum per-unit routing cost. We require the following result.

Lemma 2 *Suppose that G is a tree, and u, v, and w are distinct nodes such that $G - uv + uw$ is a tree. Then:*

$$C_u(G) = C_u(G - uv + uw).$$

In other words, the cost to u remains the same in both graphs.

Proof of Lemma. The proof follows by three basic facts. First, since the same number of edges are incident on u in G and $G - uv + uw$, the link maintenance cost for u remains the same. Second, since u is connected to all nodes, it does not experience any disconnectivity cost in G or $G + uv - uw$. Finally, it is straightforward to check that the same flow crosses u in both G and $G - uv + uw$, so all costs to u are the same in both graphs. ∎

By the preceding lemma, $C_u(G) = C_u(G - uv + uw)$. Since $c_v \leq c_w$, the traffic matrix is uniform all-to-all, and v and w are in the same connected component of $G - uv$, it follows that $\Delta C_v(G - uv, uv) \leq \Delta C_w(G - uv, uw)$. Monotonicity now implies that $Q(u, v; G) \leq Q(u, w; G - uv + uw)$, and $Q(u, v; G) = Q(u, w; G - uv + uw)$ if and only if $\Delta C_v(G - uv, uv) = \Delta C_w(G - uv, uw)$. Thus if breaking the link uv and adding the contract (u, w) makes node w strictly better off, it must make node u strictly worse off, as required. ∎

The preceding proposition shows that although inefficient pairwise stable equilibria exist, any tree where only minimum routing cost nodes appear in the interior is a pairwise stable equilibrium. This is of critical importance: in particular, any star centered at a node u with $c_u \leq c_v$ for all v can thus be sustained as a pairwise stable equilibrium. It is not difficult to establish that among all forests, such a star has the lowest social cost, i.e., the highest efficiency. (See [3] for details.) In particular, we obtain the conclusion that *the most efficient minimally connected topology is a pairwise stable equilibrium.* We will establish in Section 6 that myopic dynamics *always* converge to a topology of the form assumed in the preceding proposition. Thus the proposed dynamics select a "good" equilibrium from the set of pairwise stable equilibria.

5 Dynamics

This section describes *myopic best response dynamics* for our network formation game. Myopic dynamics refer to the property of the dynamics that at any given round, nodes update their strategic decisions only to optimize their current payoff. This is in contrast to dynamics that consider some long-term objective. We have two complementary desiderata for the proposed dynamics. First, since the Nash equilibrium is an inadequate solution concept we focus on pairwise stability. Thus we would like our dynamics to converge to a pairwise stable equilibrium. Our second objective involves efficiency: we aim to ensure that such dynamics lead to *desirable* pairwise stable equilibria. The remainder of our paper presents conditions on the contracting function that ensure that the myopic dynamics converge and the desiderata are satisfied.

We consider a discrete-time myopic dynamics that includes two stages at every round. At round k, both a node u_k and an edge $u_k v_k$ are *activated*. At the first stage of the round, with probability $p_d \in [0, 1]$, node u_k can choose to unilaterally break the edge $u_k v_k$ if it is profitable to do so; and, with probability $1 - p_d$, the link (and thus all contracts associated with) $u_k v_k$ is broken, regardless of node u_k's preference. In the second stage, u_k selects a node w and proposes to form the contract (u_k, w) to w, with associated payment given by the contracting function. (Although the second stage appears to be a restricted form of bilateral deviation, we will show that in the cost model we consider, it is sufficient to only consider bilateral deviations this form.) Node w then decides whether to accept or reject, and the dynamics then continues to the next round given the new triple of network topology, contracting graph, and payment matrix. It is crucial to note that u_k's strategic decisions are made so that its utility is maximized *at the end of the round*. We contrast this with w's strategic decision, which is made to maximize its utility at the end of the second stage *given* its utility at the end of the first stage.

We consider two variations on our basic model of dynamics: either $p_d = 1$, or $p_d < 1$. When $p_d = 1$, node u_k can choose to break either or both of the contracts associated with $u_k v_k$ (if they exist). When

$p_d < 1$, provided all links are activated infinitely often, all links are broken infinitely often *regardless of the activated node's best interest*. For ease of exposition, unless otherwise stated, all the subsequent discussion assumes $p_d = 1$.

More formally, we call an *activation process* any discrete-time stochastic process $\{(u_k, v_k)\}_{k\in\mathbb{N}}$ where the pairs (u_k, v_k) are i.i.d. random pairs of distinct nodes from V drawn with full support[3]. A realization of an activation process is called an *activation sequence*.[4]

The next example considers a natural activation process.

Example 2 (Uniform Activation Process) The activation process is said to be *uniform* if, for all k, u and v, $u \neq v$, the probability that $(u_k, v_k) = (u, v)$ is uniform over all ordered pairs. Thus

$$\mathbb{P}\left[(u_k, v_k) = (u, v)\right] = \frac{1}{n(n-1)}$$

Let (u_k, v_k) be the pair selected at the beginning of round k. Let $\left(G^{(k)}, \Gamma^{(k)}, \mathbf{P}^{(k)}\right)$ be the state at the beginning of the round. In a single round k, our dynamics consist of two sequential stages, as follows:

1. *Stage 1:* If $u_k v_k \in G^{(k)}$, then node u_k decides whether to break the contract (u_k, v_k) (if it exists), the contract (v_k, u_k) (if it exists), or both.

2. *Stage 2:* Node u_k decides if it wishes to form a contract with another w_k. If it chooses to do so, then u_k asks to form the contract (u_k, w_k), and w_k can accept or reject. The contract is added to the contracting graph if w_k accepts the contract.

Node u_k takes actions in stages 1 and 2 that maximize its utility *at the end of the round*; in the event no action can strictly improve node u_k's utility in a stage, we assume that u_k takes no action at that stage. Note, in particular, that at stage 1 node u_k only breaks (u_k, v_k) and/or (v_k, u_k) if a profitable deviation is anticipated to be possible at stage 2. At stage 2, node w_k accepts u_k's offer if this yields a higher utility to w_k than the state *at the beginning of stage 2*.

At stage 2, we only consider a very specific bilateral deviation between u_k and w_k; this is consistent with our discussion above. Further, at stage 2, there may be multiple choices w_k that maximize the utility of node u_k; while in principle this may yield nondeterministic dynamics, we defer discussion of this issue to the conclusion of the section.

[3]This implies that every pair (u, v) is selected infinitely often.

[4]We note that all the results in this paper can be proved under the following generalization of an activation process. Let u, v, w and x be four nodes from V such that $u \neq v$ and $w \neq x$. We can define an activation process to be any sequence of pairs of nodes such that, almost surely, all two pairs of nodes (u, v) and (w, x) are activated successively infinitely often.

The rules for updating the contracting graph $\Gamma^{(k+1)}$, at the end of round k, are summarized in Table 1. The first three actions described in table 1 are the basic actions the first node of the selected pair can do during a round. The last two actions are compositions of two of the basic actions.

We define $G^{(k+1)}$ to be the associated network topology: i.e., $ij \in G^{(k+1)}$ if and only if either $(i,j) \in \Gamma^{(k+1)}$ or $(j,i) \in \Gamma^{(k+1)}$ (or both). In all cases, the payment vector $\mathbf{P}^{(k+1)}$ is updated as in (2), first after stage 1, and then after stage 2.

Table 1: Updating the contracting graph

Action(s) selected by u_k	$\Gamma^{(k+1)}$
Breaks (u_k, v_k)	$\Gamma^{(k)} \setminus \{(u_k, v_k)\}$
Breaks (v_k, u_k)	$\Gamma^{(k)} \setminus \{(v_k, u_k)\}$
Adds (u_k, w_k)	$\Gamma^{(k)} \bigcup \{(u_k, w_k)\}$
Breaks (u_k, v_k) and (v_k, u_k)	$\Gamma^{(k)} \setminus \{(u_k, v_k), (v_k, u_k)\}$
Breaks (u_k, v_k) and adds (u_k, w_k)	$\left(\Gamma^{(k)} \setminus \{(u_k, v_k)\}\right) \bigcup \{(u_k, w_k)\}$

Observe that the *state* of the dynamics at round k, $\left(G^{(k)}, \Gamma^{(k)}, \mathbf{P}^{(k)}\right)$, need not be a *feasible outcome* in general. This follows because the payment matrix may not be consistent with the current contracting graph: when contracts are updated, only payments associated to the added or deleted contracts are updated—all other payments remain the same (cf. (2)). This motivates the following definition which will be used when discussing feasibility of outcomes.

Definition 2 (Adaptedness) *Let $(G, \Gamma, \mathbf{P})$ be a triple consisting of a (undirected) network topology, a (directed) contracting graph, and a payment matrix. We say that the edge ij is* adapted *in $(G, \Gamma, \mathbf{P})$ if the following conditions hold:*

1. *If $(i,j) \in \Gamma$, then $P_{ij} = Q(i,j;G)$; otherwise $P_{ij} = 0$.*
2. *If $(j,i) \in \Gamma$, then $P_{ji} = Q(j,i;G)$; otherwise $P_{ji} = 0$.*
3. *$ij \in G$ if and only if either $(i,j) \in \Gamma$ or $(j,i) \in \Gamma$.*

Note that if every edge ij is adapted to $(G, \Gamma, \mathbf{P})$, then $(G, \Gamma, \mathbf{P})$ must be a feasible outcome. Further, note that if the initial state of our dynamics was a feasible outcome, then condition 3 of the preceding definition is satisfied in every round.

The following definition captures convergence.

Definition 3 (Convergence) *Given any initial feasible outcome* $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P}^{(0)}\right)$ *and an instance of the activation process* AP*, we say the dynamics* converge *if there exists* K *such that, for* $k > K$

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

Further, we say that the dynamics converge uniformly if for every $\varepsilon > 0$ *there is* K *such that*

$$\Pr\left[\left(G^{(k+1)}, \Gamma^{(k+1)}, \mathbf{P}^{(k+1)}\right) = \left(G^{(k)}, \Gamma^{(k)}, \mathbf{P}^{(k)}\right) \ \forall k \geq K\right] \geq 1 - \varepsilon,$$

where the probability is taken w.r.t. the AP.[5]

(We say that *the network topology converges* if the preceding condition is only satisfied by G^k while $\Gamma^{(k)}$ and $\mathbf{P}^{(k)}$ keep changing, and similarly for uniform convergence of the topology.) We emphasize that the convergence we consider is uniform over realizations of the activation process; in particular, this easily implies almost sure convergence, via an application of the Borel-Cantelli lemma. Note that in our definition of convergence, we do not require that the payments between nodes in the limiting state have any relation to the contracting function; this will be established in the convergence results.

As noted above, the active node at a round, say u, may not have a unique utility-maximizing choice of a "partner" node at stage 2. To avoid oscillations induced by the possibility of multiple optimal choices, we introduce the following assumption of *inertia*. Let u_k be the node activated at round k, and suppose that at the start of stage 2 in round k, u_k has multiple utility-maximizing choices of nodes w_k. Then we assume that among such utility-maximizing nodes, u_k chooses the node w_k *it was connected to most recently*, or at random if no such node exists; *this assumption remains in force throughout the paper*. 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_k, if node u_k always chooses the node w_k with the highest degree, our convergence results remain valid.

The dynamics we have defined address an inherent tension. On one hand, any dynamic process must allow sufficient exploration of bilateral deviations to have any hope of converging to a pairwise stable equilibrium. On the other hand, if the dynamics are completely unconstrained—for example, if nodes can choose any bilateral or unilateral deviation they wish—then we have little hope of converging to an efficient pairwise stable equilibrium. Our dynamics are designed to allow sufficient exploration without sacrificing efficiency, under reasonable assumptions on the contracting function and the cost model.

[5] Note that the limiting state is a random variable, since its value depends on the activation sequence realized.

6 Convergence Analysis

In this section we prove that, under an anti-symmetric and monotone contracting function, the dynamics previously defined converge to a pairwise stable outcome where the network topology is a tree, and where non-leaf nodes have minimum per-unit routing cost. In the special case where there exists a unique minimum per-unit routing cost node $u_{\min}$, our result implies that the dynamics always converge to a star centered at $u_{\min}$. Note that other, less efficient pairwise stable outcomes may exist; thus in this special case, our dynamics converge to a feasible outcome that minimizes the price of stability. Further, we prove that, if $p_d < 1$ (i.e. if all links are broken exogenously infinitely often), then the results still hold even when the contracting function is only monotone. In all that follows let $V_{\min} = \{i \in V : c_i \leq c_j \text{ for all } j \in V\}$. Thus $V_{\min}$ is the set of all nodes with minimum per-unit routing cost.

We begin by relating the cost model of (3) to the dynamics proposed in Section 5. We will assume that Assumption 1 holds; as a result, as suggested by Corollary 1 and Proposition 2, we can expect two implications. First, nodes will break links until the graph is minimally connected. Second, if the graph is minimally connected at the beginning of a round, then it must remain so at the end of the round; thus, if u_k's action breaks the link $u_k v_k$ at the first stage of round k, then the bilateral deviation at the second stage must involve formation of exactly one link. Note that this observation serves as justification of the bilateral deviation considered at stage 2 of our dynamics for, at the second stage, we need only to consider deviations where u_k either identifies a node w_k with which to establish the contract (u_k, w_k), or does nothing.[6]

The following theorems are the central results of this paper. Our first result establishes convergence of our dynamics when the contracting function is anti-symmetric and monotone, and $p_d = 1$.

Theorem 2 *Suppose that Assumption 1 holds, and that the contracting function is monotone and anti-symmetric. Let* $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ *be a feasible outcome. Then for any activation process, the dynamics initiated at* $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ *converge uniformly. Further, if the activation process is a uniform activation process, then the expected number of rounds to convergence is* $O(n^5)$*. For a given activation sequence, let the limiting state be* $(G, \Gamma, \mathbf{P})$*. Then:*

1. *G is a tree where any node that is not a leaf is in* $V_{\min}$.
2. $(G, \Gamma, \mathbf{P})$ *is a pairwise stable outcome.*

As the proof is somewhat lengthy, we defer it to Appendix A, and only sketch the proof here.

[6] In general, the directionality of the contract may affect the payment; however, in the case of anti-symmetric contracting functions, whether (u_k, w_k) or (w_k, u_k) is formed will not impact the payment made across the contract.

Proof sketch. The proof proceeds in three main steps (for the uniform activation case).

1. *Convergence to a tree.* We first show that the network topology converges to a tree. More precisely we show that in expectation, after $O(n^4)$ rounds:

   (a) $G^{(k)}$ is a tree; and

   (b) If (u, v) and (v, u) are both in $\Gamma^{(k)}$, then $P_{uv}^{(k)} = P_{vu}^{(k)} = 0$.

2. *Convergence of the network topology.* Next, we show that the network topology converges. In particular, we show that in expectation, after an additional $O(n^5)$ rounds, the network topology is a tree where all non-leaf nodes are in $V_{\min}$. Further, the network topology remains constant in subsequent rounds.

3. *Convergence of the contracting graph.* The remainder of the proof establishes that the contracting graph converges: in expectation, after an additional $O(n^3)$ rounds, the contracting graph remains constant, and all edges are adapted (and remain so). □

When $p_d < 1$, we obtain an even stronger result regarding dynamics: we can prove that monotonicity of the contracting function suffices to establish convergence; anti-symmetry is no longer required.

Theorem 3 *Suppose Assumption 1 holds, and that the contracting function is monotone. Further, assume that $p_d < 1$. Let $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ be a feasible outcome. Then for any activation process, the dynamics initiated at $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ are such that the network topology converges uniformly.*

For a given activation sequence, let the limiting network topology be G. Also, let K be such that, $G^k = G$ for all $k > K$. Then, for $k > K$ sufficiently large :

1. *G is a tree where any node that is not a leaf is in $V_{\min}$.*
2. *$(G, \Gamma^k, \mathbf{P}^k)$ is a pairwise stable outcome.*

The proof of this second theorem requires some mild modifications to the proof of Theorem 2; in Appendix B, we point out where we explicitly use $p_d < 1$ instead of anti-symmetry. It is important to note that, if the contracting function is not anti-symmetric, convergence of the network topology does not imply convergence of the contracting graph. Nevertheless, our result is surprising as it states that, although the contracting graph might not converge, the network topology always converges. Further, after a finite time, all outcomes exhibited are pairwise stable. In Appendix B we show that if p_d is inversely polynomial in n, then the expected time to convergence is polynomial as well.

The following corollary addresses an important special case; it follows immediately from Theorems 2 and 3.

Corollary 4 *Suppose Assumption 1 holds and the contracting function is monotone. Suppose in addition that either: (1)* $p_d = 1$ *and the contracting function is anti-symmetric; or (2) that* $p_d < 1$*. Suppose in addition that* $V_{\min}$ *consists of only a single node* $u_{\min}$*. Given* $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ *and an activation sequence, let* $(G, \Gamma, \mathbf{P})$ *be the limiting pairwise stable outcome. Then* G *is a star centered at* $u_{\min}$ *and is therefore efficient.*

The preceding results demonstrate the power of the dynamics we have defined, as well as the importance of the assumptions made on the contracting functions. Despite the fact that our model may have many pairwise stable equilibria, our dynamics select "good" network topologies as their limit points *regardless of the initial state*. At the very least, only nodes with minimum per-unit routing cost are responsible for forwarding traffic (cf. Theorems 2 and 3); and at best, when only a single node has minimum per-unit routing cost, our dynamics select the network topology that minimizes social cost among all forests. This result suggests that from a regulatory or design perspective, monotone anti-symmetric contracting functions have significant efficiency benefits.

7 Conclusion

There are several natural open directions suggested by this paper. The most obvious one is to expand the strategy space considered by each node in our dynamics. More precisely, it would be interesting to analyze the robustness of the results when the active node *can select which link to break* during phase 1. A second obvious direction is to consider activation sequences that are not independent of the network state. Though our proofs rely on each link being broken infinitely often, the results can probably be extended to the case where such a property is not de-facto assumed. Finally, while our model is entirely heterogeneous in the assumptions made about the routing costs of nodes, we require the traffic matrix to be uniform all-to-all, and all links to have the same formation cost π. We intend to study the extension of the model defined here to heterogenous formation cost and traffic matrix.

Our work opens up the opportunity to consider the formation aspect for ad-hoc network design protocols. While we focused on pairwise stability and efficiency, there are other measures that might be of interest to the network designer such as fairness or resilience to link failure. As long as the dynamics satisfies the conditions stated above, the convergence to a stable outcome is guaranteed. Restricting the dynamics more could lead to additional beneficial properties.

8 Acknowledgments

The authors benefited from helpful conversations with John N. Tsitsiklis and Eric Dallal. This work was supported in part by the National Science Foundation under grant CMMI-0620811, and by the Media X Program at Stanford University and by the Canada Research Chairs Program.

References

[1] E. Altman, T. Boulogne, R. El-Azouzi, T. Jimnez, and L. Wynter. A survey on networking games in telecommunications. *Computers & Operations Research*, 33(2):286–311, 2006.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation, 2004.

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

[4] K. Arrow. Rationality of self and others in an economic system. *The Journal of Buisiness*, **59**(4):S385–S399, 1986.

[5] V. Bala and S. Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1230, September 2000. available at http://ideas.repec.org/a/ecm/emetrp/v68y2000i5p1181-1230.html.

[6] B. Briscoe, V. Darlagiannis, O. Heckman, H. Oliver, V. Siris, D. Songhurst, and B. Stiller. A market managed multiservice Internet (M3I). *Computer Communications*, 26(4):404–414, 2003.

[7] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, United Kingdom, 2004.

[8] P. Dubey. Inefficiency of Nash equilibria. *Mathematics of Operations Research*, 11:1–8, 1986.

[9] S. Eidenbenz, V. Kumar, and S. Zust. Equilibria in topology control games for ad hoc networks. In *Proceedings of International Conference on Mobile Computing and Networking*, pages 2–11, 2003.

[10] A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *Proceedings of the 22nd Annual ACM Symposium on Principles of Distributed Computing*, pages 347–351, 2003.

[11] M. Falkner, M. Devetsikiotis, and I. Lambadaris. An overview of pricing concepts for broadband IP networks. *IEEE Communications Surveys*, 3(2), 2000.

[12] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press, Cambridge, Massachusetts, 1998.

[13] M. O. Jackson. A survey of models of network formation: Stability and efficiency. Working Paper 1161, California Institute of Technology, Division of the Humanities and Social Sciences, Mar. 2003. Available at http://ideas.repec.org/p/clt/sswopa/1161.html.

[14] M. O. Jackson and A. Watts. The existence of pairwise stable networks. *Seoul Journal of Economics*, 14(3):299–321, 2001.

[15] M. O. Jackson and A. Watts. The evolution of social and economic networks. *Journal of Economic Theory*, 106(2):265–295, 2002.

[16] M. O. Jackson and A. Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, October 1996. available at http://ideas.repec.org/a/eee/jetheo/v71y1996i1p44-74.html.

[17] R. Johari, S. Mannor, and J. N. Tsitsiklis. A contract-based model for directed network formation. *Games and Economic Behavior*, 56(2):201–224, August 2006. available at http://ideas.repec.org/a/eee/gamebe/v56y2006i2p201-224.html.

[18] R. Komali and A. MacKenzie. Distributed topology control in ad hoc networks: A game theoretic perspective. In *Proceedings of IEEE CCNC*, pages 563–568, 2006.

[19] S. Mannor and J. Shamma. Multi-agent learning for engineers. *Artificial Intelligence*, pages 417–422, 2007. Special Issue on Foundations of Multi-Agent Learning.

[20] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, Oxford, United Kingdom, 1995.

[21] M. J. Osborne and A. Rubinstein. *Bargaining and Markets*. Academic Press, San Diego, California, 1990.

[22] T. Roughgarden. *Selfish routing and the price of anarchy*. MIT Press, Cambridge, Massachusetts, 2005.

[23] J. Shamma and G. Arslan. Dynamic fictitious play, dynamic gradient play, and distributed convergence to nash equilibria. *IEEE Transactions on Automatic Control*, pages 312–327, 2005.

A Proof of Theorem 2

In this appendix we prove Theorem 2. In our proof we explicitly assume a uniform activation process. Our proof continues to hold even if the activation process is not uniform; however, the running time estimate will no longer necessarily be $O(n^5)$, as stated in the theorem. For simplicity we omit the generalization to an arbitrary i.i.d. activation process.

Our proof has three major components. We first prove that, in expectation, the network topology is a tree after $O(n^4)$ rounds (Section A.1). We then prove that in an additional $O(n^5)$ rounds, the network topology converges (Section A.2). Finally, we prove that in an additional $O(n^3)$ rounds, the contracting graph converges as well, and all edges become adapted (Section A.3). We complete the proof in Section A.4.

In all that follows, we take the assumptions of Theorem 2 as given. We assume that Assumption 1 holds, and that $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ is a feasible outcome. Further, we assume that the contracting function is anti-symmetric and monotone. We also let $E^{(k)}$ denote the set of edges in the network topology $G^{(k)}$.

A.1 Convergence to a Tree

In this subsection, we prove that the network topology becomes a tree after $O(n^4)$ rounds, in expectation; further, once the network topology becomes a tree, it remains so.

Lemma 3 *Suppose that u_k and w are connected in the network topology after stage 1 of round k, where u_k is the active node. Then the contract (u_k, w) will not be added in stage 2 of round k.*

Proof. Let G be the network topology after stage 1 of round k. If $u_k w \in G$, then it is clear that both u_k and w cannot increase their utilities by adding the contract (u_k, w); so assume without loss of generality that $u_k w \notin G$.

Given that traffic is routed using shortest paths, adding $u_k w$ to G does not decrease the traffic routing cost incurred by u_k or w. Furthermore, u_k and w are connected, hence the disconnectivity cost is unchanged if $u_k w$ is added to G. Finally, adding $u_k w$ to G would increase the maintenance cost incurred by both nodes. Thus both nodes' cost would increase if (u_k, w) is added at stage 2. For u_k to offer the contract (u_k, w) in stage 2, u_k must receive a positive payment in return. However, in that case w would not accept the contract, as claimed. ■

Corollary 5 *Assume that the network topology $G^{(k)}$ is connected. Then $|E^{(k+1)}| \leq |E^{(k)}|$.*

Proof. The proof follows from Lemma 3: if an edge is added at stage 2 of round k, then an edge must have been deleted in stage 1, leaving the total number of edges unchanged. ■

Corollary 6 *If both (u, v) and (v, u) are in $\Gamma^{(k)}$, then $P_{uv}^{(k)} = -P_{vu}^{(k)}$.*

Proof. We prove this by induction on k. The base case $k = 0$ is true by the anti-symmetry assumption, since $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P}^{(0)}\right)$ is a feasible outcome. Assume the result is true at round k, and that it is not true at round $k + 1$. By the definition of the dynamics, we can only add at most one contract in round k. Thus there is exactly one link, say uv, for which the result is not true in round $k + 1$. But this implies that one of (u, v) or (v, u) was added during round k while the link uv was already in place. This is a contradiction to Lemma 3. ■

Lemma 4 *In expectation, after $O(n^2)$ rounds, $G^{(k)}$ is connected.*

Proof. We first prove that the number of connected components in the network topology is non-increasing. Then we show that during any round where the network topology is disconnected, with probability at least $1/n$, the number of connected components decreases. Finally, a bound on the maximum number of connected components of any given graph yields the result.

We prove the first step by contradiction. From the definition of the dynamics, if the number of connected components increases at round k, then a link is broken at stage 1 and no link between different connected components is formed at stage 2. Lemma 3 then implies that in fact, *no* contract is added in stage 2. However, Assumption 1 implies that at stage 2, there is at least one node w in a different connected component from the current active node u_k, such that both u_k and w are strictly better off forming the contract (u_k, w)—a contradiction.

To prove the second step, assume that $G^{(\ell)}$ is the network topology at the beginning of the ℓ'th round. Further, assume that $G^{(\ell)}$ has at least two connected components. Then the number of edges not in $G^{(\ell)}$ is at least $n - 1$. Thus, with probability at least $1/n$, we have $u_\ell v_\ell \notin G^{(\ell)}$, hence no link will be broken during stage 1. By Assumption 1, at stage 2 a link will be formed between u_ℓ and a node in a different connected component from u_ℓ. Thus with probability at least $1/n$, the number of connected components decreases.

Thus, in expectation, after n rounds the number of connected components decreases by one. The number of connected components in $G^{(0)}$ is at most n. Hence, in expectation, after $O(n^2)$ rounds the network topology is connected. ■

The following corollary follows from the fact that the number of connected components in the network topology is non-increasing, as shown in the first part of the proof of Lemma 4,

Corollary 7 *If $G^{(k)}$ is connected, then $G^{(k+1)}$ is connected.*

We now know that in expectation, within $O(n^2)$ rounds, the network topology is connected *regardless of the initial network topology $G^{(0)}$*. Further, once the network topology becomes connected, it remains connected. We now prove that in expectation, in another $O(n^4)$ rounds, the network topology is a tree. Further we prove that once the network topology becomes a tree, it remains a tree.

Lemma 5 *Assume there is a cycle in the network topology $G^{(k)}$. Then for each link of the cycle, at least one of the endpoints' utility would increase if the link was removed.*

Proof. Without loss of generality, let uv be a link from the cycle such that $P_{uv}^{(k)} - P_{vu}^{(k)} \geq 0$. Similar to the proof of Lemma 3, we know that $C_u\left(G^{(k)}\right) > C_u\left(G^{(k)} - uv\right)$, since u saves the link maintenance cost, the traffic routing cost can only decrease, and the disconnectivity cost remains unchanged. We conclude that u's utility would increase if the link uv was removed. ■

Lemma 6 *In expectation, after $O(n^4)$ rounds, $G^{(k)}$ is a tree. Further, if $G^{(\ell)}$ is a tree for some ℓ, then the network topology is a tree in all later rounds.*

Proof. Corollary 5 and Corollary 7 imply that if $G^{(\ell)}$ is a tree, then $G^{(k)}$ is a tree for all $k \geq \ell$. We know by Lemma 4 and Corollary 7 that, in expectation, after $O(n^2)$ rounds, $G^{(k)}$ is connected, and remains so thereafter. Thus we assume without loss of generality that $G^{(k)}$ is connected. To complete the proof, by linearity of expectation, we only need to prove that in expectation, in $O(n^4)$ additional rounds, the network topology is a tree.

We first show that if $G^{(k)}$ contains a cycle, then with probability at least $1/n^2$, $|E^{(k)}| > |E^{(k+1)}|$. From Lemma 5, we know that at least one link of the cycle, say uv, is such that u would be better off if the link uv was removed. Assume that at round k, $(u_k, v_k) = (u, v)$, which happens with probability $1/n(n-1) > 1/n^2$. From Corollary 6, we conclude that the *best* action for u in stage 1 is to remove all contracts associated with uv. At the end of stage 1, the network topology is still connected. By Lemma 3, no contract will be added in stage 2, as required. Hence, in expectation, after $O(n^2)$ rounds, the number of edges decreases by one.

At most $O(n^2)$ edges must be removed from $G^{(k)}$ to yield a tree; and as long as the network topology is connected, the number of edges cannot increase (cf. Corollary 5). Thus, we conclude that in expectation the network topology is a tree after $O(n^4)$ rounds. ■

A consequence of Lemma 6 and Corollary 6 is the following key proposition.

Proposition 4 *In expectation, after $O(n^4)$ rounds:*

1. *$G^{(k)}$ is a tree; and*
2. *If (u, v) and (v, u) are both in $\Gamma^{(k)}$, then $P_{uv}^{(k)} = P_{vu}^{(k)} = 0$.*

Further, (1) and (2) hold in all later rounds.

Proof. Lemma 6 implies the first claim; we establish the second. We know that in expectation, after $O(n^4)$ rounds, $G^{(\ell)}$ is a tree. For each of the $n-1$ links in $G^{(\ell)}$, Corollary 6 implies that if both contracts are present, then their payments have opposite sign.

Assume that $uv \in G^{(\ell)}$ is such that (u, v) and (v, u) are both in $\Gamma^{(k)}$, and $P_{uv}^{(\ell)} > 0$; then, of course, $P_{uv}^{(\ell)} = -P_{vu}^{(\ell)}$. With probability $1/n(n-1)$, at round ℓ we have $(u_\ell, v_\ell) = (u, v)$. Since $P_{uv}^{(\ell)} > 0$, the best action for u at stage 1 is to break at least one of the two contracts associated with uv. By Lemma 3, the link uv will never have two associated contracts again.

There are at most $n-1$ links with two associated nonzero contracts. Thus in expectation, after $O(n^3)$ rounds, any remaining links with two associated contracts must have zero payments associated with both contracts. ■

A.2 Convergence of the Network Topology

In this subsection, we establish that in an additional $O(n^5)$ rounds, in expectation, the network topology converges. In all that follows, our starting point is Proposition 4; effectively, it allows us to assume that the network topology is a tree at the beginning of any round after $O(n^4)$ rounds have passed. In particular, *we assume the following without loss of generality for the duration of this subsection.*

Assumption 2 *There exists a round k such that for all $\ell \geq k$, $G^{(\ell)}$ is a tree; and if (u, v) and (v, u) are in $\Gamma^{(\ell)}$, then $P_{uv}^{(\ell)} = P_{vu}^{(\ell)} = 0$.*

Since the network topology is a tree, and remains so, the network topology can only change if a link is broken in stage 1 of a round, and a new link is formed in stage 2. The following lemma characterizes this sequence of actions.

Lemma 7 *Suppose that u_k breaks the link $u_k v_k$ in stage 1 of round k. Then at stage 2, the link $u_k w_k$ is formed, where:*

$$w_k = \arg\min\{c_w : G^{(k)} - u_k v_k + u_k w \textit{ is connected}\}. \tag{4}$$

In other words, u_k connects to a node w_k that has minimum per-unit routing cost, among all nodes that yield a connected network topology.

Proof. By Assumption 2, $G^{(k)}$ is a tree, and we have already shown that $G^{(k+1)}$ will be a tree as well. Thus we know u_k will only connect to a node in stage 2 such that the resulting network topology is connected. Among all such nodes, by monotonicity of the contracting function and Lemma 2, u_k will prefer to connect to the node that has minimum per-unit routing cost. Node w_k will accept the connection because of Assumption 1. ∎

We now make use of the following potential function:

$$F(G) = \sum_{i \notin V_{min}} d_i(G). \tag{5}$$

Thus F is the sum of the degrees of nodes *not* in $V_{\min}$. Note that we only evaluate F when G is a tree, so that $n > F(G) \geq F_{\min} = n - |V_{\min}|$. The latter bound is achieved if and only if G is a tree where the only non-leaf nodes are in $V_{\min}$.

Lemma 8 *The potential function is non-increasing, i.e. $F\left(G^{(k+1)}\right) \leq F\left(G^{(k)}\right)$.*

Proof. Given the definition of F, and given that both $G^{(k)}$ and $G^{(k+1)}$ are trees, the potential function value can increase by at most one in any round. This follows since at most one edge can be deleted in stage 1, and at most one edge can be added in stage 2 and both such edges share an endpoint. Suppose that (u_k, v_k) is activated at round k, and that $F\left(G^{(k+1)}\right) = F\left(G^{(k)}\right) + 1$. Then u_k must have broken the link $u_k v_k$ in stage 1, and created a link with a node not in $V_{\min}$ in stage 2. Further, we must have $v_k \in V_{\min}$, since otherwise we have $F\left(G^{(k+1)}\right) = F\left(G^{(k)}\right)$. This contradicts Lemma 7. ∎

Lemma 9 *Suppose that $F\left(G^{(k)}\right) > F_{\min}$. Then with probability at least $1/n^4$, F strictly decreases in two rounds.*

Proof. Given that $F\left(G^{(k)}\right)$ is not minimal, there is a node $v \notin V_{min}$ of degree at least 2. Let v_1 and v_2 be two of its neighbors. Let $u \in V_{min}$. Given that $G^{(k)}$ is a tree, at least one of v_1 and v_2 is connected to u via v; assume without loss of generality this node is v_1 (i.e., v lies on the path from v_1 to u). By Assumption 2, at most one of $\{P_{vv_1}^{(k)}, P_{v_1v}^{(k)}\}$ is nonzero. We start by assuming that $P_{v_1v}^{(k)} = 0$; a symmetric analysis applies in the other case.

We split the analysis into two disjoint cases. First assume that $P_{vv_1}^{(k)} \le Q\left(v, v_1; G^{(k)}\right)$. Assume that the active pair is (v_1, v), which happens with probability $1/n(n-1) > 1/n^4$. By anti-symmetry and monotonicity, we have that:

$$
\begin{aligned}
Q\left(v_1, u; G^{(k)} - vv_1 + v_1u\right) &< Q\left(v_1, v; G^{(k)}\right) \\
&= -Q\left(v, v_1; G^{(k)}\right) \\
&\le -P_{vv_1}^{(k)}.
\end{aligned}
$$

Note that $-P_{vv_1}^{(k)}$ can be interpreted as the net payment v_1 currently makes to v. Given that $\Delta C_{v_1}(G - vv_1, v_1u) = 0$ by Lemma 2, we conclude it is strictly profitable for v_1 to break (v, v_1) in stage 1, and form a contract to some node in $V_{\min}$ in stage 2 (since the minimum in (4) is achieved by u). Thus the potential function is reduced in one round.

Instead, assume that $P_{vv_1}^{(k)} > Q\left(v, v_1; G^{(k)}\right)$. Assume that the next two activated pairs are (v, v_1) and (w^*, v), where w^* minimizes $Q\left(v, w; G - vv_1 + vw\right)$ among all nodes w such that $G - vv_1 + vw$ is connected. This sequence of activations takes place with probability $[1/n(n-1)]^2 > 1/n^4$. By reasoning similar to the preceding paragraph, in round k, v will break the link vv_1 in stage 1 and form vw^* in stage 2. If $w^* \in V_{\min}$, then the potential function decreased in one round. If $w^* \notin V_{\min}$, then in round $k+1$ the payment for the contract (v, w^*) will be $Q\left(v, w^*; G^{(k+1)}\right)$. Given that $G^{(k)}$ was connected and $w^* \notin V_{\min}$, we are in the first case we considered: we can replace v_1 by w^*, and repeat the argument of the preceding paragraph. Thus the potential function will decrease after two rounds.

A symmetric analysis can be carried out if $P_{v_1v}^{(k)} \neq 0$ instead. We conclude that the potential function decreases after two rounds with probability at least $1/n^4$. ∎

Note that when $p_d \neq 1$, then we only need to adjust the probability of the event to p_d/n^2, since with that probability the desired pair is selected and the link is broken exogenously.

The following corollary establishes convergence of the network topology.

Corollary 8 *In expectation, after $O(n^5)$ rounds, the network topology is a tree where all non-leaf nodes are in $V_{\min}$. Further, the network topology remains constant in subsequent rounds.*

Proof. From Lemmas 8 and 9, we conclude that, in expectation, after $O(n^5)$ rounds, the potential remains constant at $F_{\min}$. Thus, the network topology is a tree where all non-leaf nodes are from $V_{\min}$. By Lemma 7 and the assumption of inertia, the network topology can no longer change. ■

A.3 Convergence of the Contracting Graph

We now only need to prove that once the network topology has converged, the contracting graph will also converge. In this subsection we establish that in an additional $O(n^3)$ rounds, in expectation, the contracting graph converges. *We assume the following without loss of generality for the duration of this subsection.*

Assumption 3 *There exists a round k such that the network topology has converged to a tree G where all non-leaf nodes are in $V_{\min}$; i.e., for all $\ell \geq k$, $G^{(\ell)} = G$. Further, for all $\ell \geq k$, if (u, v) and (v, u) are in $\Gamma^{(\ell)}$, then $P_{uv}^{(\ell)} = P_{vu}^{(\ell)} = 0$.*

Lemma 10 *In expectation, after $O(n^3)$ rounds, the contracting graph remains constant, and all edges are adapted (and remain so).*

Proof. Under Assumption 3, it suffices to show that, in expectation, all links are adapted after $O(n^3)$ rounds. By definition of our dynamics, any edges $uv \notin G$ will have $P_{uv}^{(\ell)} = 0$ for all $\ell \geq k$. Thus we restrict attention to edges $uv \in G$.

Note that the number of non-adapted edges cannot increase in any round after round k, since the network topology remains constant. Suppose that edge $uv \in G$ is not adapted in round ℓ. We prove that with probability at least $1/n^2$, the number of non-adapted links decreases by one in round ℓ. Since uv is not adapted, we can assume without loss of generality that $(u, v) \in \Gamma^{(k)}$ and by Assumption 3, $P_{uv}^{(k)} - P_{vu}^{(k)} > Q\left(u, v; G^{(k)}\right)$. With probability $1/n(n-1) > 1/n^2$, $(u_\ell, v_\ell) = (u, v)$; i.e., (u, v) is activated at round ℓ. In this case, u's best sequence of actions is to break uv in stage 1, and re-establish the contract (u, v) in stage 2. Thus, at the conclusion of round ℓ, the link uv will be adapted, so the number of non-adapted edges decreased by one in round ℓ.

There are at most $n-1$ non-adapted links in $G^{(k)}$. Thus, in expectation, all links in $G^{(k)}$ become adapted (and remain so) after $O(n^3)$ rounds. ■

A.4 Completing the Proof

We are now ready to prove Theorem 2. We restate it for completeness.

Theorem 2. *Suppose Assumption 1 holds, and that the contracting function is monotone and anti-symmetric. Let* $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ *be a feasible outcome. Then for any activation process, the dynamics initiated at* $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ *converge uniformly. Further, if the activation process is a uniform activation process, then the expected number of rounds to convergence is* $O(n^5)$.

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

1. *G is a tree where any node that is not a leaf is in $V_{\min}$.*
2. *$(G, \Gamma, \mathbf{P})$ is a pairwise stable outcome.*

Proof. The expected number of rounds until the network topology becomes a tree is $O(n^4)$ (by Proposition 4). The expected number of additional rounds until the network topology converges is $O(n^5)$ (by Corollary 8). From that point, the expected number of additional rounds until the contracting graph and payment matrix converge is $O(n^3)$ (by Lemma 10). Thus the expected number of rounds to convergence is $O(n^5)$.

Since all edges are adapted in the limiting payment matrix (by Lemma 10), the limiting state is a feasible outcome. The network topology is a tree where any non-leaf node is in $V_{\min}$ (by Corollary 8). By Proposition 3, this feasible outcome is pairwise stable. ■

B Bound on Expected Convergence Time for $p_d < 1$

In all that follows, we assume that p_d is such that $p_d < 1$, and $1/p_d = O(n^a)$ for some constant $a > 0$; thus we assume that p_d is inversely polynomial in n. We can now prove the following modification of Theorem 3.

Theorem 9 *Suppose Assumption 1 holds, and that the contracting function is monotone. Further, assume that p_d is inversely polynomial in n. Let* $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ *be a feasible outcome. Then for any activation process, the dynamics initiated at* $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ *are such that the network topology converges uniformly. Further, if the activation process is a uniform activation process, then the expected number of rounds to convergence is polynomial in n.*

For a given activation sequence, let the limiting network topology be G. Also, let K be such that, $G^k = G$ for all $k > K$. Then, for $k > K$ sufficiently large :

1. *G is a tree where any node that is not a leaf is in $V_{\min}$.*
2. *$(G, \Gamma^k, \mathbf{P}^k)$ is a pairwise stable outcome.*

Proof. The proof is structured as that of Theorem 2 and we only provide technical details where necessary.

Given that $p_d < 1$, the expected number of rounds before a given link is broken is $O(\frac{1}{p_d}n^2) = O(n^{2+a})$. There are at most $O(n^2)$ links in $G^{(0)}$, thus, in expectation, after at most $O(n^{4+a})$ rounds all links from $G^{(0)}$ are broken at least once. From the definition of our dynamics, at most one contract is added at each round. Further, Lemma 3 still holds, thus after a polynomial number of rounds, every link in the network have exactly one contract associated to it. This is a stronger result than that of Corollary 6. All other results from Section A.1 now hold, and thus, in expectation, after a polynomial number of rounds, the network topology is a tree.

The proof from all the results from Section A.2 hold except from that of Lemma 9. We give here the correct statement and proof of such lemma.

Lemma 11 *Suppose that $F\left(G^{(k)}\right) > F_{\min}$. Then with probability at least p_d/n^2, F strictly decreases in one round. Thus F strictly decreases after polynomially many rounds.*

Proof. Given that $F\left(G^{(k)}\right)$ is not minimal, there is a node $v \notin V_{min}$ of degree at least 2. Let v_1 and v_2 be two of its neighbors. Let $u \in V_{min}$. Given that $G^{(k)}$ is a tree, at least one of v_1 and v_2 is connected to u via v; assume without loss of generality this node is v_1 (i.e., v lies on the path from v_1 to u).

With probability $1/n^2$, the pair (v_1, v) is activated. Further, independently, with probability p_d, the link v_1v is exogenously broken. Thus, with probability at least p_d/n^2, the link v_1v is broken after stage 1. Given that v was in the path from v_1 to u and that the network topology was a tree, after stage 1 u and v_1 are in different connected components. Thus v_1 will add a link to u (or another node in $V_{\min}$) during stage 2. ■

The rest of the proof is identical to that of Theorem 2, which completes the proof of this theorem. ■