Network Formation: Bilateral Contracting and Myopic Dynamics

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. We consider a network formation game where a finite number of 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 best strategy, nodes optimize their single period payoff only. We characterize a simple set of assumptions under which these dynamics will converge to a pairwise stable network topology; 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. Further, we show that these assumptions are satisfied by a contractual model motivated by bilateral Rubinstein bargaining with infinitely patient players.

1 Introduction

Given the reliance of modern society on data networks, it is remarkable that such networks — particularly the Internet — are in fact "networks of networks". They are held together through a federation of independently owned and operated service providers, that compete and cooperate to provide service. If we wish to understand how the network will evolve under decisions made by independent self-interested network operators, then we must turn our attention to the strategic analysis of *network formation games* (NFGs).

NFGs describe the interaction between a collection of nodes that wish to form a graph. Such models have been introduced and studied in the economics literature; see, e.g., [1, 2, 3]. We consider a game theoretic model where each of the nodes in the network is a different player, and a network is formed through interaction between the players. We are interested in understanding and characterizing the networks that result when individuals 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 eventual network structure. As a specific example, we are motivated by the interaction between Internet service providers (ISPs) to form connections that yield the fabric of the global Internet. Most contractual relationships between ISPs may be classified into one of two types: *transit*, and *peer*. Provider A provides transit service to provider B if B pays A to carry traffic originating within B and destined elsewhere in the Internet (either inside or outside A's network). In such an agreement, provider A accepts the responsibility of carrying any traffic entering from B across their interconnection link. In peering agreements, one or more bidirectional links are established between two providers A and B. In contrast to transit service, where traffic is accepted regardless of the destination, in a peering relationship provider B will only accept traffic from A that is destined for points *within* B, and vice versa. (For details on Internet contracting, see [4, 5, 6].)

We highlight several key points about the contracting between ISPs that motivates the high level questions addressed in this paper. First, notice that although any given end-to-end path in the Internet may involve multiple ISPs, the network is connected only thanks to *bilateral contracts* between the different providers. Second, the ISPs use (by and large) a relatively limited set of contracts in forming connections with each other. At a high level, this motivates an important question: what contracting structures are likely to lead to "good" network topologies?

In this paper, we study this high level question through a particular network formation model. We assume that each node in the network represents a selfish agent. Motivated by data networks where links are physically present, we assume that each node participating in a link incurs a fixed maintenance cost per link. We further assume that every node is interested in sending traffic to every other node. Thus we assume that they incur a disconnectivity cost per unit of traffic they do not successfully transmit, and that nodes' experience a per-unit routing cost when forwarding or terminating traffic.

We assume that a link in the network is formed as the result of a *contract* the two nodes participating in the link at some point agreed upon. It is natural to assume that such contract induces a transfer of utility between the two nodes, and that the amount transferred is a function of the topology of the network *when the contract was formed*. We view contracting from a design perspective: what types of contracts lead to good eventual outcomes? To abstract this notion, we define a *contracting function*. If two nodes decide to form a link in a given network topology, the contracting function gives the value of the contract: both direction and amount of payment between the nodes. Given our cost structure and this notion of contracting, a given network topology together with the associated set of contracts defines the utility of an agent in the network.

Given the NFGs, we define pairwise stability of an outcome in the spirit of [3]. However, networks are not static objects; agents might negotiate a contract at a given time, but that contract might become unattractive as the structure of the network evolves. We consider dynamics that account for bilateral deviations of nodes that are assumed to be selfish and myopic. The main questions our paper answers are the following: under what conditions *on the contracting function* do the dynamics converge? When the dynamics converge, are the limiting networks pairwise stable? Are they Pareto efficient? We will find a remarkably simple set of conditions under which a form of myopic best response dynamics converge to efficient pairwise stable equilibria. Note that the dynamics we consider differ significantly from that of [7, 8] in that they account separately for both unilateral and bilateral deviations.

We also note that several other papers have also considered bilateral network formation games with transfers among the agents, including [9, 10, 1, 11, 12]. Our work differs from these earlier works by combining a network formation model where cost is incurred due to routing of traffic as well as link formation and maintenance, with the question of characterization of contracting functions that yield good limiting network topologies dynamically.

The remainder of the paper is organized as follows. We first define the class of network formation games considered in Section 2; in particular, we develop the notion of contracting in such games. In Section 3, we define *pairwise stable equilibrium*, and highlight the potential tension between pairwise stability and Pareto efficiency. In Section 4 we define and discuss the dynamics studied. Section 5 specializes our model to a particular case of interest: a network formation game with traffic routing. In Section 6 we establish the main convergence results for our myopic dynamics, in the network formation game with traffic routing.

2 The Game and Contracting

In this section, we present a network formation game where agents are the set of nodes of the network. Nodes receive value that depends on the network topology that arises. We model a scenario where each link in the network is the result of bilateral "contracting" between nodes. Each contract carries with it some utility transfer from the node seeking the agreement, to the node accepting it; we assume that the value of the utility transfers depends only on the network topology realized after agreement. We assume this contracting function satisfies certain natural properties.

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 our network formation game. We assume throughout that all edges in G are *undirected*; we use ij to denote an undirected edge between i and j. As all models in the paper address only a fixed set of nodes V, 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 $TU_i(\mathbf{P}) = \sum_j P_{ji} - P_{ij}$; the first term is the sum of payments received by i, while the second term is the sum of payments made by i. Thus the total utility of node i in graph G is $U_i(\mathbf{P}, G) = TU_i(\mathbf{P}) + v_i(G)$.

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 $ij \in G(\mathbf{T}, \mathbf{F})$.

In our model of network formation, we also assume that 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 refer to the *directed* graph $\Gamma(\mathbf{T}, \mathbf{F})$ as the *contracting graph*. 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.

The contracting graph and the network topology together determine the transfers between the nodes. Formally, we assume the existence of 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. Thus given the 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, 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 $(G, \Gamma, \mathbf{P})$.

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 [13]. We investigate this example in further detail in Appendix A.

We will be interested in contracting functions exhibiting two natural properties: *monotonicity* and *anti-symmetry*. We start with some additional notation: given $j \neq i$, let the cost to node i in network topology G be defined as $C_i(G) = -v_i(G)$. We 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 $\Delta C_i(G, ij) = 0$.)

We first define monotonicity.

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 $\Delta C_j(G, ij) > \Delta C_k(G, ik)$ 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.

Our second property is inspired 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 $Q(i, j; G) = -Q(j, i; G)$.

Note that in the game we are considering, 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. First, because nodes act as self-interested players, we define a reasonable game-theoretic notion of equilibrium for our model, called *pairwise stability* (first introduced by Jackson and Wolinsky [3]). 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. However, we are also interested in system-wide performance from a global perspective, and for this purpose we must 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 *Nash equilibrium*. However, as is commonly observed, Nash equilibrium lacks sufficient predictive power in many network formation games due to the presence of trivial equilibria.

The problem with Nash equilibrium is that 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 *both* unilateral and *bilateral* deviations. This notion is known as *pairwise stability*. It follows that any pairwise stable outcome is a Nash Equilibrium.

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. With this caveat in mind, we formalize our definition of pairwise stability as follows; note that it is similar in spirit to the definition of Jackson and Wolinsky [3].

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:*

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

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 typically 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 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 stable outcome.

The following lemma yields a useful property of pairwise stable outcomes; for the proof, see [14].

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.
Let $(G, \Gamma, \mathbf{P})$ and $(G', \Gamma', \mathbf{P}')$ be two feasible outcomes. 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 [15], 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 Dynamics

This section proposes a *myopic best response dynamic* for our network formation game. Myopic dynamics refer to the fact that at any given round, nodes update their strategic decisions only to optimize their current payoff. We have two complementary objectives in the dynamics we propose. First, we would like our dynamics to be consistent with the potential for *bilateral* deviations by pairs of nodes. Ultimately, our goal is to ensure that our dynamics always converge to a pairwise stable equilibrium. Our second objective involves efficiency: we aim to ensure that such dynamics lead to *desirable* pairwise stable equilibria. Note that this is a significant departure from the usual approach in the literature on learning in games (see, e.g., [16]), which is typically focused on ensuring convergence to some equilibrium without regard to efficiency. The remainder of our paper presents a simple set of conditions on the contracting function that ensure precisely the desired convergence results on the dynamics, in the case of a network formation game with routing.

Informally we consider a discrete-time myopic dynamic 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 later see that in the cost model we consider, it is sufficient to restrict to bilateral deviations this form.) Node w then decides whether to accept or reject, and play 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 will be made assuming $p_d = 1$.

This informal discussion leads to the following definitions. 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. A realization of an activation process is called an *activation sequence*. (In fact, all 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 next example considers a natural activation process.

Example 1 (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 $\mathbb{P}[(u_k, v_k) = (u, v)] = 1/(n(n-1))$.

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 in the state *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*. (Tie-breaking is discussed at the end of the section.)

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.

It is critical to 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*. 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.

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)	$\left(\Gamma^{(k)} \setminus \{(u_k, v_k)\}\right) \bigcup \{(u_k, w_k)\}$

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* (1) *if $(i, j) \in \Gamma$, then $P_{ij} = Q(i, j; G)$; otherwise $P_{ij} = 0$; if* (2) *if $(j, i) \in \Gamma$, then $P_{ji} = Q(j, i; G)$; otherwise $P_{ji} = 0$; and* (3) *$ij \in G$ if and only if $(i, j) \in \Gamma$ 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 activation process AP, we say the dynamics* converge *if, almost surely, there exists K such that, for $k > K$,*

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

For a given activation sequence and initial feasible outcome, we call the limiting state $(G, \Gamma, \mathbf{P})$.

(We say that *the network topology converges* if the preceding condition is only satisfied by $G^{(k)}$.) 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; we will establish such a connection in our 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.

We emphasize that the dynamics we have defined here 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.

The remainder of the paper formalizes the claim of the preceding paragraph, in a specific cost model motivated by network routing. We define our model in the next section, and study stability and efficiency in the context of this model. We then show in Section 6 that weak assumptions on the contracting function are sufficient to establish that the dynamics presented in this section always converge to a desirable pairwise stable equilibrium. In particular, when $p_d = 1$, we show that anti-symmetry and monotonicity of the contracting function suffice to establish convergence. If $p_d < 1$, then we do not need the assumption of anti-symmetry: monotonicity of the contracting function alone suffices to establish convergence.

5 A Traffic Routing Utility Model

In this section we define a network formation game where nodes extract some utility per unit of data they successfully send through the network, and study pairwise stability and efficiency in the context of this model. However, nodes experience per-unit routing costs when in the data network, as well as maintenance costs per adjacent link. Our motivation is the formation of networks in data communication settings, such as wireless ad hoc networks. Such networks are typically highly reconfigurable, with a tradeoff between costs for both link maintenance and disconnectivity.

We start by describing our traffic routing model. Formally, we suppose 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. Further, we assume that in case of multiple shortest paths of equal length, traffic is split equally among all available paths. We 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 $R_i(G) = c_i f_i(G)$.

We next turn our attention to network maintenance costs. We assume that each node experiences a maintenance cost $\pi > 0$ per link incident to it. Note that this maintenance cost is incurred by both endpoints of a link, so that the effective cost of a single link is 2π. Further, note that the link maintenance cost does not depend on the identities of the endpoints of the link; this homogeneity assumption is made for technical simplicity. 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.

Finally, nodes' experience a disconnection cost that is decreasing in the amount of traffic successfully sent. An equivalent way to view this cost is to assume that links receive an increasing utility in the amount traffic sent. Formally, we assume that each node experiences a cost of $\lambda > 0$ per unit of traffic not sent. Note that λ is identical for all nodes; again, this homogeneity assumption simplifies the technical development. Thus given a graph G, the cost to a node i from incomplete connectivity, or disconnection cost, is $D_i = \lambda(n - n_i(G))$, 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:

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

5.1 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 [14].

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.)

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 [14].

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 [14] 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:*

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

5.2 Efficiency of Equilibria

Pairwise stable equilibria will typically be inefficient (see [14] for explicit constructions of an arbitrarily inefficient equilibrium). If we restrict our attention to minimally connected pairwise stable equilibria, one can see that a star centered at $u_{\min}$ would generate lower social cost than any other minimally connected 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.*

The key result we require is the following.

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 the cost to u is the same in both graphs.*

The preceding proposition shows that although inefficient pairwise stable equilibria may exist, any tree where only minimum routing cost nodes appear in the interior is also sustainable as 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 [14] for details.) In particular, we obtain the important conclusion that *the most efficient minimally connected topology can be sustained as a pairwise stable equilibrium.* We will establish in Section 6 that our dynamics *always* converge to a topology of the form assumed in the preceding proposition. Thus our dynamics select a "good" equilibrium from the set of pairwise stable equilibria.

6 Convergence Results

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 4. We consider a model where λ satisfies Assumption 1; 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.[4]

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 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. 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* $V_{\min}$*; and (2)* $(G, \Gamma, \mathbf{P})$ *is a pairwise stable outcome.*

As the proof is somewhat lengthy, we only sketch it here. Details can be found in [14].

Proof sketch. The proof proceeds in three main steps.

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

[4] 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.

and, 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 converges to a tree where all non-leaf nodes are in $V_{\min}$. (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 get 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 the dynamics initiated at $\left(G^{(0)}, \Gamma^{(0)}, \mathbf{P^{(0)}}\right)$ are such that, for any activation process, the network topology converges.*

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 $V_{\min}$; and (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. 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 very 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. If p_d is inversely polynomial in n, then the expected time to convergence is polynomial as well. Details can be found in [14].

The following corollary addresses an important special case; it follows 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 the unique minimally connected efficient network topology: a star centered at $u_{\min}$.*

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. Though our proofs rely on each link being broken infinitely often, it seems natural to believe that the results can 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 such settings.

Acknowledgments

The authors benefited from helpful conversations with Eric Dallal, John N. Tsitsiklis and Matthew O. Jackson. This work was supported by DARPA under the Information Theory for Mobile Ad Hoc Networks Program, by the Media X Program at Stanford, by the Canada Research Chairs Program, by the NSERC, and by the NSF under grant CMMI-0620811.

References

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

[2] 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 (March 2003) Available at http://ideas.repec.org/p/clt/sswopa/1161.html.

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

[4] Gao, L., Rexford, J.: Stable Internet routing without global coordination. IEEE/ACM Transactions on Networking **9**(6) (2001) 681–692

[5] Norton, W.: A business case for peering. Available on request from: http://www.equinix.com/resources/whitepapers (2002)

[6] Norton, W.: Evolution of the U.S. Internet peering ecosystem. Available on request from: http://www.equinix.com/resources/whitepapers (2003)

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

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

[9] Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation (2004)
[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. (2003) 511–520
[11] Bloch, F., Jackson, M.O.: The formation of networks with transfers among players. Working Papers 2004.80, Fondazione Eni Enrico Mattei (May 2004) Available at http://ideas.repec.org/p/fem/femwpa/2004.80.html.
[12] Johari, R., Mannor, S., Tsitsiklis, J.N.: A contract-based model for directed network formation. Games and Economic Behavior **56**(2) (August 2006) 201–224 available at http://ideas.repec.org/a/eee/gamebe/v56y2006i2p201-224.html.
[13] Osborne, M.J., Rubinstein, A.: Bargaining and Markets. Academic Press, San Diego, California (1990)
[14] Arcaute, E., Johari, R., Mannor, S.: Network formation: Bilateral contracting and myopic dynamics. Technical report, Stanford University Management Science and Engineering (2007)
[15] Mas-Colell, A., Whinston, M.D., Green, J.R.: Microeconomic Theory. Oxford University Press, Oxford, United Kingdom (1995)
[16] Fudenberg, D., Levine, D.K.: The Theory of Learning in Games. MIT Press, Cambridge, Massachusetts (1998)
[17] Muthoo, A.: Bargaining Theory with Applications. Cambridge University Press, Cambridge, United Kingdom (1999)

A Rubinstein Bargaining and Contracting

In this appendix we derive the contracting functions associated to solutions of a *two player Rubinstein Bargaining game of alternating offers*. We first derive the contracting function when players are infinitely patient. This corresponds to the *cost sharing* case. We then derive the contracting function for the general case.

We begin with the *Cost Sharing Contracting Function*. The cost sharing contracting function is defined by:

$$Q(i,j;G) = \begin{cases} \frac{1}{2}\left(\Delta C_j(G-ij,ij) - \Delta C_i(G-ij,ij)\right), & \text{if } ij \in G; \\ 0, & \text{otherwise.} \end{cases}$$

The cost sharing contracting function has the property that if a link ij is added to the network topology, the resulting total change in utility to i and j is equally shared between them. Formally, suppose $ij \notin G$, and that the contract (i,j) is formed. Then the total change in the utility of node i is:

$$-\Delta C_i(G,ij) - Q(i,j;G+ij) = -\frac{1}{2}\left(\Delta C_i(G,ij) + \Delta C_j(G,ij)\right).$$

Similarly, the total change in the utility of node j is:

$$-\Delta C_j(G,ij) + Q(i,j;G+ij) = -\frac{1}{2}\left(\Delta C_i(G,ij) + \Delta C_j(G,ij)\right).$$

Thus both i and j experience the same change in utility; note that identical expressions emerge if the contract (j, i) is formed instead. We conclude *the net change in utility to* i *and* j *is identical, and independent of the direction of the contract formed.*

We now consider the general solution of the two player Rubinstein Bargaining game of alternating offers. We call the corresponding contracting function the *Bilateral Bargaining* contracting function.

Consider a graph G containing the link ij. The *bilateral bargaining* contracting function value $Q(i, j; G)$ is based on the outcome of a Rubinstein bilateral bargaining game of alternating offers (see [17] for more details), with the following properties:

1. Node i (resp, j) has discount factor $\delta_i \in [0, 1)$ (resp., $\delta_j \in [0, 1)$);
2. Node i makes the first offer in the bargaining game; and
3. The players are bargaining to split a common "pie," where the size of the pie is the total difference in cost to both players between the graph G and the graph $G - ij$ G, i.e., $\Delta C_i(G - ij, ij) + \Delta C_j(G - ij, ij)$.

Thus the players i and j are bargaining to split any increase or decrease in utility that accrues to the pair as a result of the formation of the link ij. The directionality in the contract corresponds to the fact that one players leads in the bargaining game. It is well known that this game has a unique subgame perfect equilibrium, in which node i makes the first offer and j immediately accepts.

The contracting function value $Q(i, j; G)$ corresponds to the payment i must make to j so that the total difference in the utilities of nodes i and j between the network topologies G and $G - ij$ matches the unique subgame perfect equilibrium of the game of alternating offers described above. Thus we wish to ensure that:

$$-\Delta C_i(G{-}ij; ij){-}Q(i, j; G) = -\left(\frac{1 - \delta_j}{1 - \delta_i \delta_j}\right)(\Delta C_i(G - ij, ij) + \Delta C_j(G - ij, ij)) .$$

Rearranging terms yields:

$$Q(i, j; G) = \left(\frac{1 - \delta_j}{1 - \delta_i \delta_j}\right)\Delta C_j(G - ij, ij) - \left(\frac{\delta_j - \delta_i \delta_j}{1 - \delta_i \delta_j}\right)\Delta C_i(G - ij, ij).$$

Note that if $\delta_i \to 1$ and $\delta_j \to 1$, then the preceding expression converges to the cost sharing contracting function described in the preceding example. Thus we can view cost sharing as the outcome of a Rubinstein bilateral bargaining game where players are infinitely patient.