Dynamics and Stability in Network Formation Games with Bilateral Contracts

Esteban Arcaute
Computational and Mathematical Engineering
Stanford University
arcaute@stanford.edu

Eric Dallal
Electrical and Computer Engineering
McGill University
eric.dallal@gmail.com

Ramesh Johari
Management Science and Engineering
Stanford University
rjohari@stanford.edu

Shie Mannor
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 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.**

We characterize stable networks, and study the efficiency of those networks. We also consider myopic best response dynamics in the case where links are bidirectional. Under certain assumptions, these myopic dynamics converge to a stable network; further, they naturally select an efficient equilibrium out of the set of possible equilibria.

I. Introduction

Network formation games 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.

Such network formation models are applicable to a wide range of engineering problems, such as topology formation in mobile ad hoc networks (MANETs) and contracting between Internet service providers. In MANETs, nodes must often use decentralized algorithms to agree on a topology to be used for routing. Such algorithms are essentially implementations of a distributed network formation process, and characterizing such formation games can help provide insight into the design of topology formation algorithms.

As a specific economic example, consider instead the interaction between Internet service providers (ISPs) to form connections that yield the fabric of the global Internet. We highlight several key points about the contracting between ISPs that motivates the high level questions addressed in this paper. First, 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. (For further details on Internet contracting, see, e.g., [4], [5], and [6].) At a high level, this motivates an important question: when are bilateral contracts 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.* 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 network formation game, we define pairwise stability of an outcome in the spirit of [2]. 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. Thus, the main questions studied in our paper are: what are the pairwise stable equilibria of the network formation game? When are they efficient? When do natural dynamics converge to these equilibria? For both the directed and undirected link cases, we are able to characterize the structure of equilibra, as well as their efficiency. For the model with undirected links, we have also found a remarkably simple set of conditions under which a form of myopic best response dynamics converge to efficient pairwise stable equilibria.

We note that several other papers have also considered bilateral network formation games with transfers among the agents, including [7], [1], [8], [9], [10]. 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. Further, few of these papers contain any discussion of dynamic convergence to good network topologies.

The outline of the remainder of the paper is as follows. In Section II, we describe the cost and utility model we consider, including routing costs, link maintenance costs, disconnectivity costs, and monetary transfers between nodes. In Section III-A, we describe a static game played by the nodes; we begin by focusing on the case with undirected links. We assume that each node can declare the nodes it wishes to connect *to*, and the nodes it is willing to accept connections *from*. When there is bilateral consent to form a link, a contract is formed, and the value of this contract is determined bilaterally. In Section V we present results on a simple dynamic procedure that leads to an equilibrium. Moreover, this happens in polynomial time (in the number of nodes), if the deviating node is selected at random. In Section VI we summarize some results for the case of network formation games for directed networks and compare them with the results for undirected networks. It turns out that there are some striking similarities between the two models: every stable network in the undirected link model corresponds to a stable network in the directed model, but not vice versa. However, the dynamics behave quite differently in the two models. Some conclusions are finally offered in Section VII.

II. Model

In this section we define a network formation game consisting of a set of nodes $V = \{1, ..., n\}$; we also refer to nodes as *players*. We formulate a game where nodes extract some utility per unit of data they successfully send through the network. However, nodes experience per-unit routing costs when in the data network, as well as maintenance costs per adjacent link. We fully formulate the utility model constructed from these elements.

A. Notation

We use the notation $G = (V, E)$ to denote a graph consisting of vertices V and edges E. We will study two models that differ in the directedness of the graph. In the directed model, edges are directed: if there is an edge (i, j), traffic only flows across this link in the direction from i to j. We discuss results related to this model in Section VI; further details can be found in [8]. Instead, for the moment we focus on the undirected model where all edges in G are *undirected* and traffic can flow in both directions; this is the model studied in [11]. 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. Finally, we let $d_i(G)$ be the degree of node i in G, and $n_i(G)$ the number of nodes i can reach in the graph G.

B. Cost Structure and Utility

In this section we outline the utility model for our network formation setting. Our utility model captures four components: (1) the cost of routing traffic; (2) the cost of maintaining the network; (3) the cost of incomplete connectivity in the network; and (4) the transfers of utility between nodes directly connected. The first three components constitute the cost structure of our model.

We start by describing our traffic routing model. For simplicity, 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. Our assumptions about traffic routing are captured in the following definition. (More general assumptions on both the traffic matrix and the routing model are considered by [8].)

Assumption 1 (Shortest path routing) *Given a graph G, we assume that 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:

$$R_i(G) = c_i f_i(G). \tag{1}$$

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:

$$M_i(G) = \pi d_i(G). \tag{2}$$

Next, 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:

$$D_i(G) = \lambda(n - n_i(G)). \quad (3)$$

Thus we can now define the total cost to a node i in a graph G by:

$$C_i(G) = R_i(G) + M_i(G) + D_i(G). \quad (4)$$

Let j be a node different from i. We define the *difference* in cost to node i between graph G and graph $G + ij$ as:

$$\Delta C_i(G, ij) = C_i(G + ij) - C_i(G). \quad (5)$$

Note that if $ij \in G$, then $\Delta C_i(G, ij) = 0$.

Finally, two nodes participating in a link can exchange utility in the form of a payment. The payment exchanged by two nodes can either be the result of bilateral negotiation between them, or a transfer imposed by some external mechanism. Formally, 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 of node i is:

$$T_i(\mathbf{P}) = \sum_{j=1}^{n} P_{ji} - P_{ij}. \quad (6)$$

The first term is the sum of payments received by i, while the second term is the sum of payments made by i.

We can now define the total utility of a node i in a graph G with a payment matrix $\mathbf{P}$:

Definition 1 (Total Utility) *Let G be a graph, $\mathbf{P}$ be a payment matrix and $i \in V$ be a node from the graph. The the* total utility *of i is*

$$U_i(\mathbf{P}, G) = T_i(\mathbf{P}) - C_i(G). \quad (7)$$

III. The Game and Contracting

In this section we introduce the game formally. We then define the notion of pairwise stability. We then briefly discuss the process of contracting between nodes.

A. The Game

In this subsection we present a one-shot game where the players are the set of nodes V. We assume the cost and utility model defined in Section II. The outcome of this game yields a graph structure G and the payment matrix $\mathbf{P}$, thus determining the utility of all the players.

The game we consider is a network formation game where nodes have two basic strategic decisions: 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 selects subsets $F_i \subseteq V$ and $T_i \subseteq V$. This selection is done simultaneously. The set T_i represents the set of nodes i wishes to connect to, and F_i represents the set of nodes i is willing to accept connections from. We let $\mathbf{T} = (T_i, i \in V)$ and $\mathbf{F} = (F_i, i \in V)$ denote the composite strategy vectors.

Naturally, 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 $\mathbf{F}$; this topology then determines the costs experienced by all nodes. Formally, we have:

$$j \in T_i, \ i \in F_j \implies ij \in G(\mathbf{T}, \mathbf{F}). \quad (8)$$

In our model of network formation, we assume that formation of a link also entails a *binding contract* between the two nodes. In the example above, 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). All such contracts give rise to a *directed* graph $\Gamma(\mathbf{T}, \mathbf{F})$, that we refer to as the *contracting graph*. Formally, we have:

$$j \in T_i, \ i \in F_j \implies (i, j) \in \Gamma(\mathbf{T}, \mathbf{F}). \quad (9)$$

The contracting graph captures the inherent directionality in the contract that is formed between two nodes: in our model a link is only formed if one node asks for the link, and the target of the request accepts. Note that if the links are themselves directed, then the contracting graph and the network topology coincide.

We assume further that when a contract (i, j) is formed from i to j, there is also an associated payment. For the moment we simply assume that this payment depends on the realized network topology $G(\mathbf{T}, \mathbf{F})$. 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:

$$P_{ij}(\mathbf{T}, \mathbf{F}) = \begin{cases} Q(i, j; G(\mathbf{T}, \mathbf{F})), & \text{if } (i, j) \in \Gamma(\mathbf{T}, \mathbf{F}); \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

It is now clear that 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 $G = G(\boldsymbol{T}, \boldsymbol{F})$, $\Gamma = \Gamma(\boldsymbol{T}, \boldsymbol{F})$, and $\boldsymbol{P} = \boldsymbol{P}(\boldsymbol{T}, \boldsymbol{F})$ to represent specific instantiations of the network topology, contracting graph, and payment matrix, respectively, given strategy vectors $\boldsymbol{T}$ and $\boldsymbol{F}$. We refer to a triple $(G, \Gamma, \boldsymbol{P})$ arising from strategic decisions of the nodes as a feasible outcome, as in the following definition.

Definition 2 *A triple $(G, \Gamma, \boldsymbol{P})$ consisting of an (undirected) network topology G, a (directed) contracting graph Γ, and a payment matrix $\boldsymbol{P}$ is a* feasible outcome *of the game if there exist strategy vectors $\boldsymbol{T}$ and $\boldsymbol{F}$ such that $G = G(\boldsymbol{T}, \boldsymbol{F})$, $\Gamma = \Gamma(\boldsymbol{T}, \boldsymbol{F})$, and $\boldsymbol{P} = \boldsymbol{P}(\boldsymbol{T}, \boldsymbol{F})$. In this case we say $(\boldsymbol{T}, \boldsymbol{F})$* give rise *to the feasible outcome $(G, \Gamma, \boldsymbol{P})$.*

B. Contracting

In both the undirected and directed network formation models, links are associated with binding contracts between the nodes. In this section we briefly elaborate on the interpretations of the contract value $Q(i, j; G)$ for a contract formed between nodes i and j.

As a first interpretation, we presume that the value of the contract 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; this approach is explored further in [11]. An important special case of Rubinstein bargaining is *cost sharing*, where both endpoints equally share the additional cost of the link formed. Alternatively, in a model with directed links, it is always the case that i must compensate j for additional cost if the directed link (i, j) is formed; in this case, the value $Q(i, j; G)$ can be any a payment that is individually rational for both nodes. (This approach is considered in [8].) Again, because this negotiation may depend on costs experienced by the nodes, it is possible that the contracting function will depend on the exogenous model parameters as well.

Alternatively, we might *constrain* the types of contracts that can be formed. To this end, we 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. In our model, the contracting function can also depend on exogenous parameters (e.g., π, λ, and the traffic matrix), but we suppress this dependence. In practice, we expect that dependence on such exogenous parameters should only arise indirectly. This may happen, for example, if the regulator specifies a rule for computing the contract value that may also depend on the costs experienced by the two nodes in the contract. Note that the regulator may in fact be the designer of a distributed network; in this case the "contract values" are simply designed to ensure the distributed decisions of the nodes lead to good global behavior.

As a final note, we always assume that the contracting function depends only on the nodes i and j entering into the contract, as well as the network topology G. In fact, all contracting functions we consider have the property that the value $Q(i, j; G)$ depends *only* on the edges in the network topology other than ij, i.e., on the graph $G - ij$. This is reasonable: if i and j consider forming the contract (i, j), they will typically evaluate the payment in terms of the alternative—i.e., the graph without the corresponding contract present.

C. Assumptions on Contracting in the Undirected Model

In the undirected link model, we restrict attention to contracting functions that satisfy two important properties: *monotonicity* and *anti-symmetry*.

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 $Q(i, j; G + ij) > Q(i, k; G + ik)$.*

Informally, monotonicity requires that the payment to form a link must increase as the burden of forming that link increases on the accepting node. Thus, if the contract (i, j) increases the cost to node j more than the contract (i, k) increases the cost to node k, then i should pay more for (i, j) than (i, k).

Our second property is inspired by the observation that, in general, $Q(i, j; G)$ is not related to $Q(j, i; G)$. A natural way for these contracts to be related is captured in the following property.

Property 2 (Anti-symmetry) *We say that the contracting function Qis "anti-symmetric" if, for all nodes i and j, and for all graphs G, we have:*

$$Q(i, j; G) = -Q(j, i; G).$$

An antisymmetric contracting function 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. Regardless of whether the contract (i, j) or (j, i) is formed, the direction and quantity of payment across the link ij remains the same.

IV. Stability and Efficiency

In this section we define the notion of *pairwise stability*; this is the equilibrium notion we use to study the game defined in the preceding section. 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 (in a sense made precise in the following subsection).

Section IV-B is then devoted to defining efficiency for our model, as well as characterizing the efficiency of pairwise stable equilibria. We adopt *Pareto efficiency* as our definition of efficiency. Under our assumptions, we discuss the characteristics of the most efficient feasible outcomes among all pairwise stable equilibria.

A. Pairwise Stability

In this subsection, we define pairwise stability, and characterize some important properties of pairwise stable outcomes. We conclude with some examples, and a characterization of pairwise stable outcomes under anti-symmetric and monotone contracting functions.

The simplest notion of equilibrium is *Nash equilibrium*; this concept requires that no *unilateral* deviations are profitable. However, as is commonly observed, Nash equilibrium lacks sufficient predictive power in many network formation games due to the presence of trivial equilibria. For example, it is not hard to see that $F_i = T_i = \emptyset$ is a Nash equilibria 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 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 *bilateral* deviations, known as *pairwise stability*. Informally, pairwise stability of a strategy vector requires that (1) no unilateral deviations are profitable; and (2) no two nodes can collude to improve their payoff.

Formally, suppose that the current strategy vectors are $\boldsymbol{T}$ and $\boldsymbol{F}$, and the current network topology and contract graph are $G = G(\boldsymbol{T}, \boldsymbol{F})$ and $\Gamma = \Gamma(\boldsymbol{T}, \boldsymbol{F})$ 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 [2].

Definition 3 *Assume Q is a contracting function. Given strategy vectors $\boldsymbol{T}$ and $\boldsymbol{F}$, let $G = G(\boldsymbol{T}, \boldsymbol{F})$, $\Gamma = \Gamma(\boldsymbol{T}, \boldsymbol{F})$, and $\boldsymbol{P} = \boldsymbol{P}(\boldsymbol{T}, \boldsymbol{F})$. Given strategy vectors $\boldsymbol{T}'$ and $\boldsymbol{F}'$, define $G' = G(\boldsymbol{T}', \boldsymbol{F}')$ and $\Gamma' = \Gamma(\boldsymbol{T}', \boldsymbol{F}')$. Define $\boldsymbol{P}'$ according to:*

$$P'_{k\ell} = \begin{cases} P_{k\ell}, & \text{if } (k,\ell) \in \Gamma' \text{ and } (k,\ell) \in \Gamma; \\ Q(k,\ell;G'), & \text{if } (k,\ell) \in \Gamma' \text{ and } (k,\ell) \notin \Gamma; \\ 0, & \text{otherwise}. \end{cases} \tag{11}$$

Then $(\boldsymbol{T}, \boldsymbol{F})$ is a pairwise stable equilibrium *if:*

1) *No unilateral deviation is profitable: for all i, and for all $\boldsymbol{T}'$ and $\boldsymbol{F}'$ that differ from $\boldsymbol{T}$ and $\boldsymbol{F}$ (respectively) only in the i'th components,*

$$U_i(\boldsymbol{P}, G) \geq U_i(\boldsymbol{P}', G').$$

2) *No bilateral deviation is profitable: for all pairs i and j, and for all $\boldsymbol{T}'$ and $\boldsymbol{F}'$ that differ from $\boldsymbol{T}$ and $\boldsymbol{F}$ only in the i'th and j'th components,*

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

Notice that (11) is a formalization of the discussion above. When nodes i and j deviate to the strategy vectors $\boldsymbol{T}'$ and $\boldsymbol{F}'$, 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 $\boldsymbol{P}'$. 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. This allows us to discuss stability of a topology without formalizing the payments made by nodes in the topology. Thus we have the following definition.

Definition 4 *A feasible outcome $(G, \Gamma, \boldsymbol{P})$ is a* pairwise stable outcome *if there exists a pair of strategy vectors $\boldsymbol{T}$ and $\boldsymbol{F}$ such that (1) $(\boldsymbol{T}, \boldsymbol{F})$ is a pairwise stable equilibrium; and (2) $(\boldsymbol{T}, \boldsymbol{F})$ give rise to $(G, \Gamma, \boldsymbol{P})$.*

Note that by our definition of the game, for all i and j such that $ij \in G$ we must have $P_{ij} = Q(i,j;G)$ in the preceding definition.

The following result is the first of a sequence of structural results that characterize the set of pairwise stable ourcomes; it follows easily since our cost structure does not value redundant links. Details are in [11].

Proposition 1 *Let $(G, \Gamma, \boldsymbol{P})$ 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 2 (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, the nodes i and j prefer to form the link ij over remaining disconnected.*

Note that if Q is anti-symmetric, this implies that if nodes i and j are not connected in G, then both are better off by forming the link ij.

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. Thus, to ensure that nodes always prefer connectivity, we will assume that λ is large enough so that it is never beneficial for a node to be in a disconnected graph. It is clear from our model that if all other model parameters are fixed, then a λ satisfying the preceding assumption must exist.

If Assumption 2 holds, we have the following corollary about pairwise stable outcomes; the proof is immediate.

Corollary 2 *If Assumption 2 holds, all pairwise stable outcomes are trees.*

From the previous corollary, we can now provide a very useful characterization of pairwise stable outcomes when Assumption 2 holds.

Proposition 3 *Suppose that Assumption 2 holds, and that Q is anti-symmetric and 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 $\boldsymbol{P}'$ as in* (11). *Then:*

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

Proof: Assume (G, Γ) is pairwise stable and $(i,j) \in \Gamma$ as well as $(j,i) \in \Gamma$, but that (without loss of generality) $Q(i,j;G) \neq 0$. First assume that $Q(i,j;G) > 0$. Then in any pairwise stable equilibrium giving rise to $(G, \Gamma, \boldsymbol{P})$, node i is paying node j a positive payment. If node i decides to remove the contract (i,j), given that $(j,i) \in \Gamma$, G will not change. Hence the payoff to i will increase by removing the contract (i,j) from Γ, so $(G, \Gamma, \boldsymbol{P})$ could not have been pairwise stable—a contradiction. A similar argument holds with node j instead of node i if $Q(i,j;G) < 0$. ■

Example 1 (Minimum Routing Cost Star) Suppose Assumption 2 holds. We consider the same model as in the preceding example, but now assume that $S_{\min}$ is a (undirected) star centered at $u_{\min}$. Assume that the contracting graph is $\Gamma_{\min}$, such that $\forall v \neq u_{\min}$, $(u_{\min}, v) \in \Gamma_{\min}$. Suppose the contracting function Q is monotone and anti-symmetric, and let $P_{ij} = Q(i,j;S_{\min})$ for $i,j \in S_{\min}$, and zero otherwise. We prove that $(S_{\min}, \Gamma_{\min}, \boldsymbol{P})$ is pairwise stable.

Given that λ is sufficiently large and $S_{\min}$ is a tree, no single node would benefit from a unilateral deviation. Since the contracting function is anti-symmetric, two nodes would never change a contract (i,j) into (j,i) , as the net utility would remain unchanged.

Given $v, u \neq u_{\min}$, let $\Gamma' = \Gamma_{\min} - (u_{\min}, v) + (u, v)$ be the graph obtained by removing the link $(u_{\min}, v)$ from $\Gamma_{\min}$ and adding the link (u,v); let S' be the corresponding network topology, and let $\boldsymbol{P}'$ be the corresponding payment matrix. To prove that $(S_{\min}, \Gamma_{\min}, \boldsymbol{P})$ is pairwise stable, it is sufficient to prove that $U_u(\boldsymbol{P}', S') < U_u(\boldsymbol{P}, S_{\min})$. Since u is a leaf in both S' and $S_{\min}$, we have $C_u(G) = C_u(S_{\min})$. Since the contracting function is monotone, and $c_{\min} < c$, we have $Q(u,v;S') > Q(u,v;S_{\min})$, establishing the result.

Proposition 4 *Suppose that Assumption 2 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.*

B. Efficiency

In this subsection we introduce the key notion of efficiency of a feasible outcome, i.e., *Pareto efficiency*. We also concentrate on the gap between pairwise stability and efficiency. Our model of dynamics is motivated by the relationship between those two concepts.

Let $(G, \Gamma,)$ and $(G', \Gamma', \boldsymbol{P}')$ be two feasible outcomes. We say that $(G, \Gamma, \boldsymbol{P})$ Pareto dominates $(G', \Gamma', \boldsymbol{P}')$ if all players are better off in $(G, \Gamma, \boldsymbol{P})$ than in $(G', \Gamma', \boldsymbol{P}')$, and at least one is strictly better off.

In our model, the payoffs to nodes are quasilinear, so we can express efficiency in a different way [12]. We start by defining social cost of a network topology.

Definition 5 (Social Cost) *Given a (undirected) network topology G, the* social cost *of G is $\mathcal{S}(G) = \sum_{i \in V} C_i(G)$.*

Pareto efficiency can then be equivalently formulated according to the following definition.

Definition 6 (Pareto Efficiency) *Let $(G, \Gamma, \boldsymbol{P})$ be a feasible outcome. We say that $(G, \Gamma, \boldsymbol{P})$ is* Pareto efficient *if and only if the following are satisfied:*

$$(i) \ G \in \arg\min_{g} S(g);$$

$$(ii) \ \sum_{i \in V} T_i(\boldsymbol{P}) = 0.$$

It is straightforward to check that, given the definition of $T_i(\boldsymbol{P})$, the second condition is always satisfied for any payment matrix. Since condition (i) is independent of the contracting function, we have the immediate conclusion that *efficiency depends only on the network topology G.*

It is important to note that in general, efficient outcomes need not to be minimally connected. For instance, assume that V has three nodes, and that the unique minimum per unit routing cost node is such that $c_{\min} = 4$. Assume that $\lambda = 40$ and $\pi = 1$. Then the only efficient outcome is the complete graph on $\{1, 2, 3\}$.

Since pairwise stable outcomes of our game are always minimally connected, we will focus our attention on asking whether pairwise stable equilibria are most efficient *among all minimally connected graphs*.

Definition 7 (MC Efficiency) *We say that a (undirected) network topology G is* MC efficient *if: (1) it is a forest; and (2) it achieves lower social cost than any other forest G', i.e., $S(G) \leq S(G')$ for all forests G'.*

We say that a feasible outcome $(G, \Gamma, \boldsymbol{P})$ is an MC efficient *outcome if G is MC efficient.*

The following proposition characterizes MC efficient outcomes; see [11].

Proposition 5 (MC Efficient Outcomes) *Suppose that Assumption 2 holds. Let $c_{\min} = \min_i c_i$; i.e., $c_{\min}$ is the minimum per unit routing cost. Let $(G, \Gamma, \boldsymbol{P})$ be a feasible outcome, and assume that G is a forest. Then $(G, \Gamma, \boldsymbol{P})$ is MC efficient if and only if G is a star centered at a node u such that $c_u = c_{\min}$.*

V. Dynamics

A central question concerning the possible outcomes is if there is a reasonable process that leads to such an equilibrium. Due to space constraints, we do not give a precise definition of the process here. Instead, we provide an informal description and refer the reader to [11] for the rigorous definitions.

Our dynamics consists of multiple rounds; in each round are two separate stages. An activation process selects the node i_k that takes action in each round k. Further, we assume the activation process also identifies a candidate node j_k. Our only assumption on the activation sequence is that it selects all successive pairs of nodes $(i_k, j_k, i_{k+1}, j_{k+1})$ infinitely often. (For the running time estimates, it suffices that nodes are activated independently according to a distribution with full support.)

At the first stage of round k, node i_k determines whether to remove the link $i_k j_k$, if such a link exists. In the second stage of round k, node i_k chooses a potential partner w_k with which to form a new link. If a new contract is formed, a payment is also induced based on the associated contract value given the current graph. Nodes take actions in each round to maximize their utility at the end of the round. To avoid oscillations in the dynamics we also have to assume that if several nodes offer the same utility to the selected node, it would select the node it was connected to last. This property is called "inertia". Details of both the dynamics and the proof of the theorem are in [11].

Theorem 6 (Convergence) *Suppose Assumption 2 holds, the dynamics have inertia, and that the contracting function is monotone and anti-symmetric. Suppose that the graph starts from a feasible outcome. Then:*

- i *The dynamics converge for any activation process;*
- ii *If the activation process activates nodes uniformly at random, then the expected number of rounds to convergence is $O(n^5)$; and*
- iii *the limit graph and payments is a pairwise stable outcome.*

VI. The Directed Links Model

In the directed links model of [8], when a link is established between i and j traffic can be sent from i to j and not vice versa. This implies that under the same cost model described in Section II, node j does not benefit from any links directed to it since it would suffer from additional traffic and would not benefit from additional connectivity. Node i should therefore compensate node j for the additional traffic through it. We provide a brief survey of [8] and highlight the differences and similarities between the undirected and directed models. The one shot game model for directed networks is as follows: node i may offer a payment to any other node j to establish a connection from i to j. Simultaneously, node j announces the amount of money it wishes to receive from i for establishing such a link. We note that the concept of "cost sharing" is not relevant here: if the link (i, j) exists, it is a pure liability for node j, so node j would never pay any share of the cost of the link.

As in the undirected model, the equilibrium concept of interest is the pairwise stable equilibrium, since in general Nash equilibrium does not carry very much predictive power. (For example, the empty graph is always a Nash equilibrium.) A basic question is if a particular graph can be a pairwise stable equilibrium under some payments. If a similar assumption to Assumption 2 is adopted, then it is shown in [8] that only minimally connected graphs can be stable. An obvious observation is that if an edge can be removed without affecting the connectivity of the graph, it will be removed. The reason is that if i is connected to j then by removing the link ij node i can only benefit– less traffic will pass through i and node i will not pay to j. Consequently, the only graphs that can be stable are minimally connected ones. It turns out that the converse is also true: for a high λ, every minimally connected graph is also an equilibrium for some configuration of contracts. Furthermore, there is a simple characterization of the set of the payments in equilibrium: for such graphs, there is an upper bound and a lower bound to the set of possible payments in equilibrium. An interesting observation that every bidirectional tree (i.e., a tree where there is an edge from i to j and from j to i) is a minimally connected graph, and we characterize exactly when such trees are stable. In some sense, we therefore have that every pairwise stable graph in the undirected network model can be mapped to a pairwise stable graph in the directed model. Still, the directed model contains many other stable graphs that cannot be mapped to undirected stable graphs.

The analysis in [8] also considers the efficiency loss in directed networks, i.e., the gap in efficiency between pairwise stable networks and efficient (or MC efficient) networks. It is shown that the most inefficient stable (directed) network can have a social cost n times larger than either efficient or MC efficient networks. It turns out that this is the same for the undirected network model. To see that consider the following graph: node 1 is connected to node 2, node i is connected to node $i-1$ and $i+1$ ($i < 1 < n$) and node n is connected to node $n-1$. (This is just a linear network with n nodes.) Suppose that $c_i = c$ for all i. It follows from Proposition 4 that this graph is stable. A straightforward calculation shows that the social cost of this graph behaves like $O(n^3)$, while the MC efficient graph (which is also stable) is a star centered in any of the nodes $2, 3, \dots, n-1$, and has social cost that scales as $O(n^2)$. The efficient graph is the complete graph, and also has cost that scales as $O(n^2)$. Another interesting connection between the directed and undirected model is that in both cases the MC efficient structures have a star shape (as a bidirectional tree in the directed case), with the node with lowest routing cost in the middle.

Due to the nature of the cost structure, the dynamics of directed model not as rich as the dynamics of the directed model. The key problem is that a synergy cannot between two nodes, as the origin node will always pay the destination node. A similar procedure to the one suggested in Section V will lead to a minimally connected graph, but analysis from this point forward becomes intractable, as evaluation of the benefit to deviating is difficult in general directed graphs with multiple shortest paths. The undirected model, on the other hand, has much more tractable dynamics, due to the fact that even once a tree is obtained (i.e., redundant links that are not needed for connectivity are removed), a node i may break a link with node j as long it can connect to another node k which is connected to node j.

VII. Conclusion

This paper has considered pairwise stable equilibria of network formation games that model the creation of networks under a natural cost structure for communication networks. We considered both networks where the links are undirected (traffic flows in both directions), such as in networks of ISPs, and networks where the links are unidirectional, such as certain ad-hoc networks. An important feature of our models is the bargaining process, where we assume the nodes negotiate a payment for establishing links between them. We characterized the pairwise stable equilibria in such networks and discussed the efficiency of the stable configurations as opposed to the social optimum. The dynamics for the undirected case were also considered. We have shown that a fairly natural procedure leads to a pairwise stable network.

The characterization of a simple myopic process that leads to an equilibrium is a key result in this work. Quoting Arrow [13], "The attainment of equilibrium requires a disequilibrium process." Understanding the behavior of simple dynamic procedures is essential for justifying the validity of equilibrium concepts as operating points for the network. Moreover, dynamics may be used as a way for selecting more efficient equilibria out of all possible ones.

There are plenty of directions for future research. Naturally, one can consider other deviations and more complex dynamic procedures. Of particular interest is the question of which equilibria the dynamics converge to with high probability. That is, can we characterize the set of networks to which a "typical" random dynamics would converge? Finally, while our model is entirely homogeneous in the assumptions made about the routing costs of nodes, the traffic matrix, and the formation cost π. We intend to study the extension of the model defined here to such settings.

VIII. Acknowledgements

This work was supported by DARPA under the Information Theory for Mobile Ad Hoc Networks Program, by the Media X Program at Stanford, and by the NSF under grant CMMI-0620811. We also gratefully acknowledge helpful conversations with Matthew O. Jackson, John N. Tsitsiklis.

References

[1] V. Bala and S. Goyal, "A noncooperative model of network formation," *Econometrica*, vol. 68, no. 5, pp. 1181–1230, September 2000, available at http://ideas.repec.org/a/ecm/emetrp/v68y2000i5p1181-1230.html.

[2] M. O. Jackson and A. Wolinsky, "A strategic model of social and economic networks," *Journal of Economic Theory*, vol. 71, no. 1, pp. 44–74, October 1996, available at http://ideas.repec.org/a/eee/jetheo/v71y1996i1p44-74.html.

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

[4] W. Norton, "A business case for peering," 2002, available on request from: http://www.equinix.com/resources/whitepapers.

[5] ——, "Evolution of the U.S. Internet peering ecosystem," 2003, available on request from: http://www.equinix.com/resources/whitepapers.

[6] L. Gao and J. Rexford, "Stable Internet routing without global coordination," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 681–692, 2001.

[7] F. Bloch and M. O. Jackson, "The formation of networks with transfers among players," Fondazione Eni Enrico Mattei, Working Papers 2004.80, May 2004, available at http://ideas.repec.org/p/fem/femwpa/2004.80.html.

[8] R. Johari, S. Mannor, and J. N. Tsitsiklis, "A contract-based model for directed network formation," *Games and Economic Behavior*, vol. 56, no. 2, pp. 201–224, August 2006, available at http://ideas.repec.org/a/eee/gamebe/v56y2006i2p201-224.html.

[9] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler, "Near-optimal network design with selfish agents," in *Proceedings of ACM Symposium on the Theory of Computing*, 2003, pp. 511–520.

[10] 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. [Online]. Available: citeseer.ist.psu.edu/anshelevich04price.html

[11] E. Arcaute, E. Dallal, R. Johari, and S. Mannor, "Network formation: Bilateral contracting and myopic dynamics," Stanford, Department of Managmenet Science and Engineering, Tech. Rep. 33, Feb. 2007, available at http://cgi.stanford.edu/group/msande-reports/cgi-bin/reportinfo.pl?07-02-272354-33.

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

[13] K. Arrow, "Rationality of self and others in an economic system," *The Journal of Buisiness*, vol. **59**, no. 4, pp. S385–S399, 1986.