Strategic Network Formation through Peering and Service Agreements

Elliot Anshelevich * Bruce Shepherd † Gordon Wilfong ‡

Abstract

We introduce a game theoretic model of network formation in an effort to understand the complex system of business relationships between various Internet entities (e.g., Autonomous Systems, enterprise networks, residential customers). This system is at the heart of Internet connectivity. In our model we are given a network topology of nodes and links where the nodes (modeling the various Internet entities) act as the players of the game, and links represent potential contracts. Nodes wish to satisfy their demands, which earn potential revenues, but nodes may have to pay (or be paid by) their neighbors for links incident to them. By incorporating some of the qualities of Internet business relationships, we hope that our model will have predictive value. Specifically, we assume that contracts are either customer-provider or peering contracts. As often occurs in practice, we also include a mechanism that penalizes nodes if they drop traffic emanating from one of their customers.

For a natural objective function, we prove that the price of stability is at most 2. With respect to social welfare, however, the prices of anarchy and stability can both be unbounded, leading us to consider how much we must perturb the system to obtain good stable solutions. We thus focus on the quality of Nash equilibria achievable through centralized incentives*: solutions created by an "altruistic entity" (e.g., the government) able to increase individual payouts for successfully routing a particular demand. We show that if every payout is increased by a factor of 2, then there is a Nash equilibrium as good as the original centrally defined social optimum. We also show how to find equilibria efficiently in multicast trees. Finally, we give a characterization of Nash equilibria as flows of utility with certain constraints, which helps to visualize the structure of stable solutions and provides us with useful proof techniques.*

*Department of Computer Science, Princeton University, Princeton NJ. Email: eanshele@cs.princeton.edu. This research supported in part by ITR grant 0311333, Bell Labs, and the NSF Postdoctoral Fellowship in Mathematical Sciences.

†Bell Labs, Murray Hill, NJ.

‡Bell Labs, Murray Hill, NJ.

1 Introduction

The formation of the Internet marked a step away from centrally planned and controlled networks, and a step towards networks that employ distributed traffic routing control. Nevertheless, the early Internet adopted a common standard and agreed upon metrics to make routing decisions. With the advent of competition in the 1990's, this ceased to be the case, and today the Internet is composed of tens of thousands of sub-networks called *Autonomous Systems* (AS), each under a single administrative authority with its own distinct goals in controlling the traffic entering and leaving its network. A number of other emerging communication networks also have the characteristic that a collection of domains, with varying self-interests, participate in such a multilateral sharing of network resources.

Network management becomes substantially more complex as a result of the inherent interdependence involved in multi-domain (multilateral) networks. For instance, there is limited ability to predict how changes in the network, or in business relationships between domains, affects current routings [20]. Even obtaining accurate estimates of current traffic conditions is a nontrivial challenge [7, 15]. Such immediate operational tasks in turn depend on the existence of a stable system of business relationships between AS's. Questions of whether certain AS's will peer and what type of *service level agreements* (SLAs) will be forged between AS's are critical for understanding the structural properties of the networks formed. This paper's focus is on these longer-term strategic factors affecting the formation of a stable network-of-networks.

It is natural to employ game theory to analyze the self-interested behaviour of domains, and several models have recently been proposed. These have fallen into two broad categories: models that address routing issues (e.g., [14, 37]) and those that study network creation (e.g., [1, 13, 31]). Our objectives are more aligned with the latter class. In particular, we introduce a network formation model called the *Local Contract Formation Game* (LCFG) in an effort to understand the complex system of business relationships between Internet domains such as AS's, residential customers, and enterprise networks. This system is at the heart of Internet connectivity and by incorporating some of the qualities of these relationships, we hope to capture the essence of these interactions so that our model will have some predictive value. We model three key elements of real-world business interactions between domains. The first element is that, unlike many models, we assume that monetary transfers and business relationships are strictly local. This models current practice where links arise as part of a bilateral agreement between the two endpoints. Such arrangements depend implicitly on the global structure of the networks and traffic demands, but are based only on a *local* bid-ask type contract between two neighboring domains [27, 28]. Second, we allow the links in the network to be one of two types: customer-provider or peer-peer [11, 12, 16, 21]. Third, we include a mechanism that penalizes domains if they drop traffic emanating from one of their customers. This models the fact that SLA penalties have become commonplace in contracts, especially those offered by core network providers [26]. Despite these features, our model remains simple enough to analyze.

In our model, we assume that there are no link capacity constraints. In other words, links have been sized accordingly to carry all possible traffic; a reasonable assumption given our focus on long-term effects, as opposed to brief outages due to traffic bursts. Also, we assume that traffic demands have a specified path they must follow. While this *fixed route* assumption ignores dynamic routing changes in a network, it seems a necessary first step to understanding whether a given configuration of business relationships is stable. Moreover, to understand routing behavior, we must first understand how/whether the links used by those routes would themselves form a stable configuration. The model choice is also driven and justified by our focus on stable interdomain business relationships, which are longer-term than routing decisions. As we will see, the fixed route model already contains considerable complexity. Future work may be to extend this to allow route generation and contract formation as a repeated game, where routes depend on the previous contracts, and contracts depend on the routes during the previous step.

We study stable business relationships by exploring the existence and structure of Nash equilibria for the game LCFG. Nash equilibria are the dominant solution concept in game theory, and correspond to locally optimal solutions. While many solution concepts are possible in this context, if local optimality is not satisfied, then the system is inherently unstable since some node (player) could act to increase their payoff. Our goals are to understand the creation of such stable networks, and how to induce good stable solutions when none would otherwise form. Interestingly, we will see there are several attributes of Nash equilibria in LCFG that match common practice in the Internet. One such is that there always exist equilibria where no money is exchanged between peers. Another is that even without any assumption on the structure of the underlying contract graph, one may assume that no node forwards traffic from one of its providers to another. This mirrors the practice of filtering out route announcements in the Border Gateway Protocol.

Local Contract Formation Game In LCFG we are given a network topology of nodes and links where the nodes (modeling domains) act as the players of the game, and the links represent potential connections or contracts that could be made. We are also given a set of demands each of which is a path in the graph. For example, in Figure 1(a), there is a demand d_1 along the path between u_1 and u_2. A demand d_i is *active* (connected) if appropriate contracts are formed on all links in the path of d_i (i.e., the links are "activated"). If d_i is active, the utility of d_i's endpoints increases (say by some value λ_i). For every active demand going through a node v, however, v suffers 1 unit of disutility, representing the fact that it costs money to transit traffic through one's servers and subdomains. We call this cost the *transit cost*. Figure 1(b) shows the change in the nodes' utility because demands d_1 and d_3 are active. The goal of each player v in this game, then, is to satisfy as many of the demands ending at v as possible, while having as few demands as possible going through v.

In such a framework without utility transfer, demands would never be activated (i.e., connected), since players have no interest in incurring transiting costs for other players. When bilateral contracts are introduced, however, two players may activate a link between them and thus commit to carrying traffic offered across the link. In fact, two types of contracts are possible: a *customer-provider contract* where one player (the customer) pays the other player (the provider) to activate the link and transit the customer's traffic; and a *peering* contract where no payment is exchanged because it is in both players' interests to activate the link. Figure 1(c) shows the payments c_i that occur because of the formation of these contracts, as well as the impact of these payments on the players' utilities.

In addition, suppose a provider accepts payment from a customer (i.e., a contract is formed between them). The customer agrees to this payment since it then expects the provider to form the necessary contracts with its other neighbors to activate demands from the customer (or the customer's customers etc). If the provider fails to make such a payment for a link, then the provider must pay a "penalty" to its customer. Figure 1(c) shows u_6 paying a penalty p to u_5, since u_6 is the provider of u_5 and u_6 failed to form a contract with u_7 that would have activated demand d_2.

Thus a player's total utility consists of the utility for having demands activated for which it is an endpoint, payments to or from neighbors for contracts formed, transit costs for active demands that go through it and finally, penalties paid

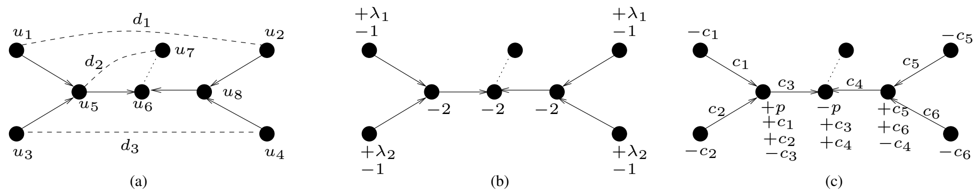

Figure 1. Utilities and transfers. Directed edges represent customer-provider contracts, with the provider at the head. The dashed edge (u_6, u_7) represents an inactive link/contract.

by or to it. A node's strategy space then consists of choosing which contracts to form and how much to ask/offer for the formation of such contracts. A node might change its strategy if by doing so it strictly increased its utility. For example, in Figure 1, if node u_6 changed its strategy and paid some amount c to node u_7 to activate the link between them, then u_6's utility would increase by p since it would no longer pay a penalty, but it would decrease by $c + 1$ due to the payment of c plus it would incur a transit cost of 1 for the now active demand d_2. Thus if $p - c - 1 > 0$, then u_6 would benefit by changing its strategy.

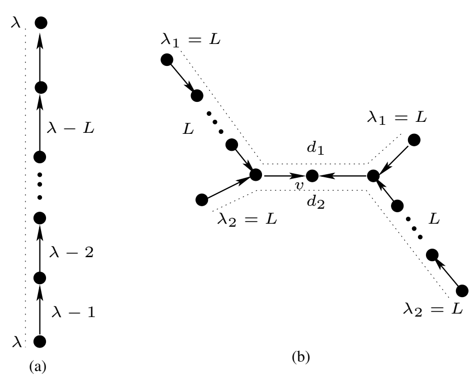

Figure 2. (a) The payments for customer-provider contracts are shown. (b) Cooperation results in a good NE.

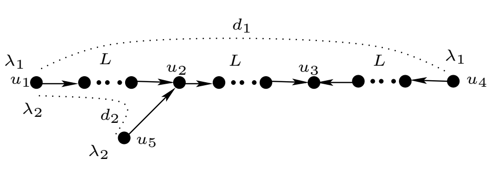

Figure 3. NE with distant monetary effects.

Nash Equilibria (NE) Formal definitions of LCFG and its NE's are given in Section 2, but here we informally illustrate some interesting forms of NE's. Notice that in any NE each player must have non-negative total utility for otherwise it could improve its utility to 0 by requesting a large ("infinite") amount for each of its incident edges effectively deactivating or "cutting" these edges. Also, a trivial NE where no demands are activated can always be achieved by having all nodes cut all of their incident edges. In this way, all nodes end up with a total utility of 0.

In Figure 2(a), we have a single demand, of utility λ, through $L + 2$ nodes. One may argue that this is a NE if $L \leq \lambda - 1$. Informally, each internal node would keep (at least) 1 from its incoming payment to pay for its transit cost, before passing on any remaining payment to the provider above it. If each node keeps exactly 1, then the payments along edges are shown in the figure. Assuming the asking prices of the providers do not exceed this amount, then the edges will be formed and the demand will be active as long as $L \leq \lambda - 1$. Note that the total utility for each internal node (contract payment received minus contract payment paid minus 1 for transit) would be 0. If a node "cuts" its outgoing edge, then it will have to pay a large penalty (in this case the penalty would actually be at least $\lambda - 1$) which is at least as large as the payment it is receiving. Thus the configuration in Figure 2(a) is stable if $L \leq \lambda - 1$.

In a sense, for this example there is sufficient utility from the bottom node of the demand to activate all of the edges. This suggests a likely candidate for a stable solution: one where each demand pays for its own transit. That is, a flow of payments is propagated from the endpoints along the demand path in order to ensure that the internal nodes are willing to incur the transit costs. There are NE's, however, as in Figure 2(b), where demands pay for each other's transit and improve the overall utility as a result. In this example, demand d_2 has plenty of money to pay for its transit on the left of node v, but requires help to pay for transit on the right, with the opposite being true for demand d_1. Both demand paths are of length $L + 4$.

More surprisingly, there are NE's where payments from a demand end up at nodes distant from that demand's connection path, as in Figure 3. Here λ_1, λ_2 are approximately L, and it can be shown that a NE exists with all the edges as active contracts. In such a NE, all the payments starting at u_1 towards u_2 would be used to pay for the transit of the two demands through those nodes from u_1 up to u_2. There is a payment from u_5 to u_2 of $\lambda_2 - 1$ which is roughly L. Thus there is about L utility available at u_5 to pay for the transit of d_1 along the path from u_2 to u_3. It is beneficial for u_2 to send the money from d_2 towards u_3 to pay for d_1's transit (i.e., pay for the provider edge from u_2 to the next node) since otherwise it would have to pay a large penalty to its customer node along the path from u_1. It would not have to pay this penalty if it also cut the edge from that customer, but then it would have to pay a penalty to u_5 since that would disconnect d_2. These examples show that paying for the transit of others can result in a NE that maximizes social welfare.[1]

Our Results As we will show, the prices of anarchy and stability[2] with respect to social welfare in LCFG can both be unbounded. Moreover, the best (maximum social welfare) NE is not poly-time approximable. This motivates the questions of when good Nash equilibria exist, and how to *induce* a good NE in an instance of LCFG given a small budget, without altering the game mechanism. Some of our results are as follows.

- In Section 3 we focus on measuring the degree to which an instance must be perturbed in order to obtain a Nash equilibrium as good as the centrally defined social optimum. Equivalently, one may view this as seeking high quality NE's achieved through *centralized incentives*; that is, solutions created by an altruistic entity able to increase individual payouts for particular demands. This entity may be the government, or any organization interested in the quality of the overall network. Our main result of this section shows that if every demand value $\lambda(d)$ is increased by a factor of 2, then there is a NE as good as the original centrally defined social optimum. We also prove generalizations of this result, that often give us factors much better than 2.
- Instead of the social welfare objective, in Section 4 we consider a non-mixed sign objective function: the total transit cost incurred plus the value of the demands that are not connected. With respect to this (still natural) objective function, we show that the price of stability of LCFG is at most 2, as well as how to efficiently find a good Nash equilibrium starting with an approximation to the centralized optimum.
- In Section 5 we show that if the routing network is a multicast tree (all demands have a node in common), then we can find the best NE in polynomial time. Unlike our other results, this problem becomes NP-hard if demands are not unit-size. For non-unit demands, we can still efficiently find an interesting non-trivial NE in a multicast arborescence.
- Finally, in Section 6, we provide a structural characterization of Nash equilibria that is also the basis for most of our technical results. We prove that every NE can be represented by a flow of utility with certain constraints. Apart from helping to visualize the general structure of stable solutions, this yields a poly-time algorithm for determining the existence of Nash equilibria that successfully route a specific set of demands.

Related Work The notion that an AS benefits from its demands getting to their destinations, and loses utility from transiting traffic, was explored in several papers (e.g. [14, 31, 35]). Typically, however, the concentration has been on short-term routing and pricing schemes, e.g., [23, 25, 30, 33, 36]. While short-term pricing and admission models are of significant interest, they rely on an underlying set of business (longer-term economic) relationships. Moreover, agreements between entities (ISPs, enterprise or residential customers) tend to be based on more rigid contracts such as fixed bandwidth, or peak bandwidth contracts [27, 28]. This is largely due to the complexity (and expense) of monitoring IP traffic at a packet or flow level.

Motivated by this, several game theoretical models have addressed the strongly related notion of network formation. Some of these do not look at contract formation, but instead assume that edges have intrinsic costs [1, 2, 4, 5, 13, 24, 34]. On the other hand, contract formation models of networks have been heavily addressed as well, mostly in the economics literature (e.g., [17], for a survey see [29]). This body of work mostly addresses questions distinct from those studied here. In particular, none of these consider customer-provider and peering contracts, or measure the impact of provider penalties. In [6] a very general model of network formation is considered. Our model is a (very) special case of theirs, and indeed their characterization of "stable networks" is analogous to our characterization of a stable set of strategies in LCFG. The model of [31] is also quite relevant; their flavor of results is quite different from ours, however, as they focus on solutions where all demands are satisfied, and on so-called pairwise-stable equilibria. In addition, [8, 30] address intra-domain concerns of when it is wise to form peering contracts, instead of concerns about the resulting overall network structure.

2 Model and Basic Results

We now formally define an instance of the Local Contract Formation Game (LCFG). We are given a mixed graph $G = (V, E)$, as described in the Introduction. Undirected edges represent possible peering contracts. A directed edge (arc) $e = (u, v)$ is referred to as a *provider edge* of u and a *customer edge* of v. Graph G could be a multi-graph, that is, there may be several types of contracts possible between two given nodes. Notice that the direction of an edge is not meant to indicate traffic flow, but to represent the hierarchical nature of a typical customer/provider network (cf. [16]).

[1] Recall that the social welfare of a solution is the sum of the players' utilities.

[2] See [32] for a definition of price of anarchy (known there as the coordination ratio), and [2] for a definition of price of stability.

We are also given a set of *demands* D on G, each specified with an unordered pair of nodes st. A demand d is a request for traffic to be exchanged between s and t; to limit notation, we assume each demand is unit size, although in most cases extending our results to demands of variable sizes is easy.

Traffic for a demand d is carried on some fixed, prespecified path $P(d)$ with endpoints s, t. A path from s to t in G is said to be an *upward path* if it is a simple directed path. We call a simple path P between s and t a *valid path* if P can be written as the concatenation of three subpaths $P = P_1 P_2 P_3$ where P_1 is a (possibly empty) upward path from s to some node u, P_2 is an empty path or a peering edge between u and some node v and P_3 is a (possibly empty) upward path from t to v. This definition of a valid path is the same as the definition of Type-1 and Type-2 paths in [16], valid signaling paths in [21] and valley-free paths in [11]. For the remainder of this paper, we assume that all our paths $P(d)$ are valid. See the end of this section for why this assumption is the correct one to make.

Notice that in the above definition $P(d)$ specifies exactly which contracts must be active for demand d to be satisfied. In fact, all our results still hold if we relax this condition, and say that demand d is satisfied when there exists any valid path from s to t on the nodes of $P(d)$ (in other words, the route is fixed, but not the contracts along this route).

A *configuration* of our game is determined by a set S of *active edges* (i.e. successfully formed contracts). A demand d is *active* (i.e., successfully routed) in this configuration if all edges of $P(d)$ are active. Let $D^{end}(v)$ be the active demands having endpoint v, and $D(v)$ be the set of active demands for which $v \in P(d)$. Similarly, for each edge e, define $D^{end}(e)$ to be the active demands having e as an initial or final edge, and $D(e)$ to be the set of active d with $e \in P(d)$. In general, for any $S \subseteq V \cup E$, $D(S)$ is the set of active demands that include nodes/edges of S. We denote by $\lambda_s(d)$ the *value* of the demand d to s if d is active and s is an endpoint of d. For any node v, the set of active edges determines a total value to v of: $\sum_{d \in D^{end}(v)} \lambda_v(d)$. In the following, we always assume that $\lambda_v(d) \geq 2$.

The basic goal of nodes (players) is to influence the formation of contracts so as to balance their desire for the values of connected demands with the cost incurred by transiting the traffic of other nodes. We now discuss the basic strategies (actions) that players have at their disposal to achieve their goals.

Strategies: A *strategy* for a node v consists of *bids* for each of its incident edges e. We denote by $\text{offer}_v(e)$ the amount v is offering (or demanding if $\text{offer}_v(e) < 0$) to form the business relationship represented by e. Similar to [6], a contract on edge e with endpoints u and v is *formed* if $\text{offer}_v(e) + \text{offer}_u(e) \geq 0$, the payment offered by one endpoint of e is greater than the payment demanded by the other endpoint. In this case e is *active* for these strategies. This is slightly different from [6], since if one of our bids is negative, then the other endpoint transfers an amount equal to its absolute value. If both endpoint bids are nonnegative, then the edge is formed with no payments between u and v. We let $c_u(e)$ denote the ultimate payment made to u for e. A *solution* $\mathcal{S}$ for this game consists of a profile of strategies, i.e., for each node a list of bids for its incident edges. Clearly $\mathcal{S}$ induces a network configuration S given by the set of active edges.

Transit Costs: For any configuration of active edges, say induced by a solution $\mathcal{S}$, define the *traffic transited* by a node v (or similarly by an edge e) to be $D(v)$ (or $D(e)$). For $x \in V \cup E$, let $t(x)$ be the total amount of traffic that x is transiting, i.e. $t(x) = |D(x)|$. We define the *cost of transiting* for a node x to be $t(x)$. That is, there is a normalized cost of 1 for each active demand transited. In general, for any $S \subseteq V \cup E$, let $t(S) = \sum_{s \in S} t(s)$.

Penalties: Finally, a provider must pay penalties to its customers if it fails to meet its obligations. The existence of such penalties is the only difference between peering and customer-provider edges in our model. For node v and demand d with $v \in P(d)$, define $p_v(d) = 0$ except in the following cases. Consider a demand d with endpoints s and t. The demand is *penalty-enabled* for v if there is a nonempty active directed subpath of $P(d)$ from an endpoint of $P(d)$ up to v, the last edge obviously being a customer-provider edge $e = (u, v)$. Assume the endpoint of $P(d)$ in this subpath is s. If v fails to form a contract activating the next edge of $P(d)$ for a penalty-enabled demand d, then v pays a penalty $p_v(d) = \lambda_s(d) - 1$ to u, and symmetrically $p_u(d) = -p_v(d)$. This models the fact that a customer of v wants to send traffic on $P(d)$, but v is unable to activate its incident edges in $P(d)$, thereby failing in its provider duties.

Note that our results hold with other penalty models as well, such as penalizing a provider v for lost demands, even if the "culprit" edges are not incident to v. Intuitively, a node s should be compensated for loss of any λ-value from one of its demands d routed through one of its (paid for) connections to a provider. The benefit to s is $\lambda_s(d) - 1$ (since s receives $\lambda_s(d)$ value if d is active but has a cost of 1 for transit of d). If penalties were any smaller, there would be instability in the system, since providers would prefer to pay a penalty instead of forwarding a customer's traffic. On the other hand, our results, including hardness results, still hold if the penalties were allowed to be higher.

The Utility Function: Given the utility for an active demand, the transiting costs, and the penalties, the utility of node v (for a solution $\mathcal{S}$) is *utility*(v) (or *utility*$_{\mathcal{S}}(v)$), as follows.

$$utility(v) = \sum_{d \in D^{end}(v)} \lambda_v(d) - t(v) + \sum_e c_v(e) - \sum_d p_v(d) \quad (1)$$

Equation (1) may seem complicated, but its components are quite intuitive. A node v gains the value of $\lambda_v(d)$ for each active demand that it originates, loses 1 for every demand it transits, gets payment $c_v(e)$ according to the contract it makes with its neighbor on e (either positive if v is paid or negative if it pays) and loses $p_v(d)$ for penalties (either positive or negative depending on whether it pays or receives the penalty).

Nash Equilibria Given a solution $\mathcal{S}$, a *deviation* for a node v is a solution where v changes its strategy while all others remain as in $\mathcal{S}$. A *best deviation* for a node v is a deviation for v that results in the highest payoff to v over all possible deviations for v. In the case that a best deviation for v is to stay with its original strategy, we say that v is *stable*. If all nodes are stable in a solution $\mathcal{S}$, then $\mathcal{S}$ is said to be a *Nash equilibrium* (or NE). We call a NE *nontrivial* if it has at least one active demand (and trivial otherwise). We say that S, a set of edges, *induces* a NE when there is a NE whose set of active edges is S.

Discussion of Filtering and Valid Paths Here we discuss the "no-filtering" and valid path assumptions in our model. An important property in our model is that strategies of a node consist of the amount of money it is offering/demanding for various connection agreements (i.e., edges). It is not part of a node's strategy to decide which demands it will transit; it must transit all active demands that pass through it. At first glance, this might seem unrealistic, since AS's can do anything they desire with traffic, including not forwarding it or filtering away particular packets. However, consider the case where we drop these restrictions, allow demands to follow paths that are not valid and allow arbitrary filtering of traffic. When would it be in the interest of node v to not transit an active demand $d \in D(v)$? If v is an endpoint of d, then v does not lose anything by transiting d, since we assume $\lambda_v(d) \geq 2$. If $d = st$ does not originate at v, then there must be two edges of $P(d)$ incident to v. If at least one of these edges is a customer edge, then v would have to pay a penalty of $\lambda_s(d) - 1$ (or $\lambda_t(d) - 1$) to its customer for not transiting d. But then the penalty would be at least as much as the cost to transit so there is no gain in filtering such traffic. In fact, the only time when v would gain by filtering a demand d is when the two edges of d incident to v are both non-customer edges. In this case, v can refuse to transit d, save itself the transit cost, and not lose any utility since it has no customers that it would owe penalties to. In fact, this is exactly the type of demand (route) filtering that is done in the Internet today [12, 16]. Because it is *always* in v's interest to filter in such a case, we can simply assume that all demands with such routes have been filtered out, which is equivalent to assuming that all demands follow valid paths and that additional filtering is unnecessary.

2.1 Basic Results and Useful Observations

In our model a trivial Nash equilibrium always exists. A solution where all players v set all offer$_v(e)$ to a large negative number results in no active edges, and is a Nash equilibrium with all nodes having a utility of 0. Moreover, for every Nash equilibrium there is an equivalent one where the payments demanded on inactive edges are infinite. Without loss of generality, we assume from this point on that this holds for all inactive edges in any stable solution we consider. Thus we can now think of deviations as a node "cutting" edges since forming extra contracts is not an option for a single player in such a solution. Cutting is achieved either by a customer offering less money or a provider requiring more money.

Intuitively (and in practice) money is paid from customers to providers (not the other way around) and peering connections typically involve no money changing hands at all. In our model definition, we did not enforce this to be so. However, this property is inherent in our definition of the players' utilities, as the following proposition illustrates (all proofs not given immediately can be found in [3], the full version of this paper).

Proposition 2.1 *For every Nash equilibrium A, there exists another Nash equilibrium B with the same active edges, and all strictly positive payments only being paid from customer to provider.*

Thus we now only consider NE's where positive payments are never made to a node's peers or customers. To simplify notation, we let $c(e)$ be the nonnegative payment on e from the customer to the provider.

Price of Anarchy and Stability An *optimal centralized solution*, that we denote by OPT, is a configuration that maximizes the social welfare function $\sum_{v \in V}$ *utility*(v). Notice that all the contract payments and penalties cancel out, so this objective function is just $\sum_{d=st \in D(V)}(\lambda_s(d) + \lambda_t(d) - \|P(d)\|)$ where $\|P(d)\|$ is the number of nodes in $\|P(d)\|$. We say that a solution is a *best Nash equilibrium* if it maximizes the social welfare function over all possible Nash equilibria. We use the notation $W(S)$ to denote the social welfare of a solution S.

The example in Figure 2(a) with $L = \lambda$ shows that there are instances of LCFG in which a best Nash equilibrium has no active demands and hence 0 social welfare, whereas OPT has non-zero social welfare. This implies that the price of anarchy [32] and price of stability [2] can both be infinite in LCFG. In this example, all the edges are active in OPT, so

$W(OPT) = 2L - (L+2) = L - 2$. However, there is no nontrivial NE in this instance. To see this, note that in a NE, each node must have non-negative utility. In order to cover the cost of transit, each internal node must keep at least 1 from any payment from their customer. This cannot happen since there are $L + 1$ such nodes, and the total payment offered by the bottom node is at most L.

Our focus now is to study when nontrivial Nash equilibria exist, and how good they can be in terms of social welfare. Unfortunately, the following holds.

Theorem 2.2 *Finding a nontrivial Nash equilibrium is NP-complete even when G is an arborescence. Moreover, there is no polynomial-factor approximation algorithm for finding the best Nash equilibrium.*

See [3] for the proof of this and other hardness results. Theorem 2.2 shows the intractability of finding NE's of any value, let alone close to $W(OPT)$. There may also not exist a good quality approximate NE (where players only deviate if they substantially improve their utility) because of the example in Figure 2(a). This drives our focus on how to add incentives to achieve a good NE.

3 Creating Good Nash Equilibria

As noted above, since all Nash equilibria in LCFG may be of very poor quality, we must allow incentive schemes if we hope to form good Nash equilibria. We consider incentives in the form of payments by some central authority to players, under the constraint of some total budget B.

There are several forms of incentives an entity could offer to players, four of which are as follows:

1. For some or all players, give some amount of money to a player only if it follows a particular strategy.
2. For some or all players, for some or all edges incident to a player, give some amount of money to a player only if it does not cut this edge.
3. Increase the penalties for not delivering particular demands.
4. Increase the values of particular demands ($\lambda_s(d)$'s).

We define a *Type-i Nash equilibrium* as a NE in an instance of LCFG where we also employ incentives of form i, $i = 1, 2, 3, 4$, as described above. In this paper, we consider Type-4 Nash equilibria, since they are more general than the others because of the following theorem.

Theorem 3.1 *Let N_i be the collection of sets of edges that induce a Type-i Nash equilibrium, with a fixed incentive budget B. Then, $N_1 = N_2 \subseteq N_4$, and $N_3 \subseteq N_4$.*

To prove this theorem, we first need to prove the useful Lemma 3.2. The proofs of both Lemma 3.2 and Theorem 3.1 are in the full version [3].

Lemma 3.2 *Suppose there is a node v with 2 best deviations, cutting S_1 and S_2. Then, cutting $S_1 \cap S_2$ is also a best deviation of v.*

3.1 A Nash Equilibrium as good as OPT

The main result of this section is that if we increase the λ-values for every demand by a factor of 2, then in the resulting game instance there is a Nash equilibrium whose active edges are exactly the active edges of OPT. For the results in this section, we assume $\lambda_s(d) = \lambda_t(d)$ for s, t being endpoints of d and we write $\lambda(d)$ to be this common value. If this were not the case, then instead of increasing $\lambda_s(d)$ by a factor of 2, the results hold if we set $\lambda_s(d)$ to $\lambda_s(d) + \lambda_t(d)$, which is still a factor of 2 increase in total.

Theorem 3.3 *Let E^* be the set of active edges in OPT. If we increase $\lambda(d)$ by a factor of 2 for every d, then E^* induces a Nash equilibrium.*

To prove this, we first need to show a sufficient condition for a given set S of active edges to induce a Nash equilibrium. Consider the following bipartite b-matching problem $MP(S)$ with node sets A and B where A has a node for every (s, d) pair where s is an endpoint of active demand d, and B contains a node for every active non-peering edge in S. For each $e \in B$ we define the *capacity* of e to be $x(e) = |D(e)| - |D^{end}(e)|$. The capacity $x(s, d)$ for each node $(s, d) \in A$ is defined as $\lambda_s(d) - 2$. For every $(s, d) \in A$, there is an edge in $MP(S)$ between (s, d) and all nodes in B representing non-peering edges in $P(d)$ directed away from s (i.e., that are reachable from s via a directed subpath of $P(d)$). The basic idea here is that if in an x-matching an amount y is matched between $s \in A$ and $e \in B$ then y utility is somehow transferred from s to the head v of e to cover some of v's transit costs.

Lemma 3.4 *If an x-matching exists in $MP(S)$, then S induces a Nash equilibrium.*

Note that the existence of a NE on S does not always imply the existence of an x-matching in $MP(S)$. Only a subset of NE's can be described by such matchings, as they are essentially solutions where the λ utility of each demand is used to pay for transit on its own demand path (although it may be used to pay for transit of other demands on those nodes). In such a solution, utility flows from a demand endpoint s, down its demand path, until it ends up at the node v that $MP(S)$ assigns to it, and is used to pay for v's transit. The NE's in Figure 2 are both of this type, while the NE in Figure 3 is not. In fact, all NE's can be interpreted as a flow of utility, as we illustrate in Section 6.

Using Lemma 3.4, we can now prove Theorem 3.3. Since we are dealing with the optimal centralized solution, cutting any set of edges in OPT decreases the social welfare. We can use this fact to form a matching $MP(OPT)$ as above after increasing all $\lambda(d)$ values by a factor of 2. For the full proof, see the full version [3].

Extensions of Theorem 3.3 We now show that there is no need to increase λ by a factor of 2, when we can instead increase it by an additive term. Let P_1 and P_2 be the two maximal directed paths in $P(d)$ for demand d starting at the two endpoints of d. Note that P_i might just consist of an endpoint of d. Let δ_i be such that $\|P(d)\| = \delta_i \|P_i\|$, where $\|P\|$ is the number of nodes in path P. These values δ_1 and δ_2 represent the imbalance of this path. If we let δ_d be the smaller of these, we can now present the following theorem.

Theorem 3.5 *To form a Nash equilibrium on the same edges as OPT, it is enough to increase $\lambda(d)$ to become $A\lambda(d) + (1 - \frac{A}{2})(\|P(d)\|/\delta_d)$, for any $0 \le A \le 2$. In particular, it is enough to set $\lambda(d)$ to be $\frac{3}{2}\lambda(d) + \frac{\|P(d)\|}{4}$.*

This theorem shows that in many graphs, it is not necessary to increase every single λ by a factor of 2. The amount of money we need actually depends on the length of demand paths, as well as the imbalance between the two directed parts of the paths. Unfortunately, assuming that $\|P_1(d)\| = \|P_2(d)\|$ does not improve these bounds, even in the special case where all the paths $P(d)$ are confluent. See [3] for more results of this flavor.

Relative Difference We can also generalize the result of Theorem 3.3 as follows. Define the *relative difference* of a Nash equilibrium N compared to OPT in terms of social welfare as $rel(N) = (W(OPT) - W(N))/\Lambda$, with $\Lambda = \sum_{d \in D(OPT)} \lambda(d)$. This says that in terms of social welfare, the difference between OPT and N is small relative to the total value of the active demands in OPT.

The following theorem gives us a bound on relative difference. For the case where $\varepsilon = 1$, this is exactly Theorem 3.3, where we can form a NE with no relative difference.

Theorem 3.6 *If for all d we set $\lambda(d)$ to $(1+\varepsilon)\lambda(d)$, then there is a NE N such that $rel(N) \le (1-\varepsilon)$.*

4 A Different Objective Function and Price of Stability

In LCFG, the social welfare can be positive or negative, and so is a mixed-sign objective. We have shown that if we consider OPT to be the solution maximizing social welfare, then the prices of anarchy and stability can both be unbounded. This often occurs with mixed-sign objectives in optimization problems, but such objectives can often be transformed into natural same-sign objective functions that give better approximation ratios (see e.g. [19]). Consider the objective of $\sum_{d \in D(V)} \|P(d)\| + \sum_{d=st \notin D(V)} (\lambda_s(d) + \lambda_t(d))$. In this objective, we want to minimize the transit cost in the entire solution, and we also look to minimize the total λ-value of demands that are *not* connected. This is the objective function used in [31]. Notice that the minimum of this objective is also the solution with maximum social welfare, so the change of objective only matters for approximations.

As in the previous section, assume that $\lambda_s(d) = \lambda_t(d)$, and call this value $\lambda(d)$. Then the following holds (there is a corresponding result for the case that $\lambda_s(d) \neq \lambda_t(d)$).

Theorem 4.1 *With respect to the non-mixed sign objective function above, the price of stability is at most 2.*

The proof of the above theorem (which can be found in the full version [3]) does not only give us an existence result, but also an approximation algorithm that runs in polynomial time. Given any solution S, we can find a Nash equilibrium that is at most twice as expensive as S. This is done by attempting to find a matching similar to $MP(S)$ from Lemma 3.4. If we find such a matching, then we are done, and otherwise obtain a set of edges that can be taken out from the current solution without greatly increasing the cost.

5 Multicast Settings

In this section, we consider the special case of multicast demands. By *multicast* we mean that all demands have one endpoint in common. Approximating the best NE in a multicast DAG is NP-hard, but in the case where the underlying graph is a tree, we can find a best Nash equilibrium in polynomial time via a complex dynamic program.

Theorem 5.1 *In a multicast tree, we can find the best Nash equilibrium in polynomial time if all demands are unit-size.*

While all the other results in this paper extend to non-unit demands, this one does not. If, however, the underlying tree is actually an arborescence (i.e. all edges are directed towards the sink), then we can find a non-trivial Nash equilibrium with nice properties, even for non-unit demands. See [3] for details.

6 Nash Equilibrium As A Flow of Payments

In this section we provide a structural characterization of Nash equilibria that helps visualize the "movement" of money in a Nash equilibrium. It is also a useful tool for many of our proofs.

In the following, we assume that S is a fixed set of edges and we look for node strategies (i.e., values $c(e)$) that induce a NE with S as the active set of edges. The basic framework is as follows. We define a flow problem in a graph $G(S)$ such that every feasible flow in $G(S)$ corresponds to a stable set of payments $c(e)$. Such "Nash flows" are obtained by stitching together n circulation problems, one local to each node. For any fixed set of edges S, our results yield a compact LP formulation for the polyhedron of Nash equilibria (viewed as payments on edges).

Before formally defining the node circulation problems, we give some intuition on the constraints for a stable set of payments. Specifically, for every node v, think of each edge incident to v as a separate entity that has a budget equal to the utility it is bringing (or taking) from v. It may possibly use its budget to pay for its own transit or to lend to other edges. We show that a critical property of any NE is that edges cannot lend too much. We call this the *bounded lending constraint.* The following remark is an informal statement of the results in this section.

Remark 6.1 *If for all v, its incident edges can pay their transit costs in v, and no edge exceeds the* bounded lending constraint *(defined below), then the payments form a Nash equilibrium. Conversely, every Nash equilibrium can be expressed in this manner.*

Nash circulations and flows. We start by forming a circulation problem at a given node v. The problem instance for v has a node v_e for every active edge e incident to v, as well as an extra node s that represents the "outside world" (see Figure 4 where 4 copies of the same node s are displayed). The node s is the origin of all utility obtained by v, as well as the sink of all utility v must spend on transit costs, provider payments, and the utility v keeps as profit.

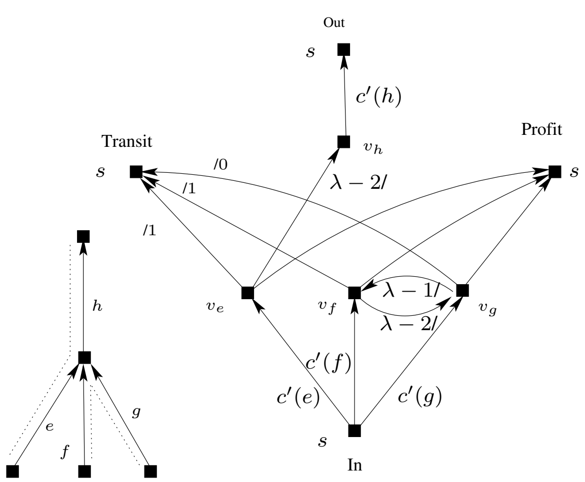

Figure 4. (Left) Node v with incident customer and provider edges. (Right) The corresponding circulation gadget, with the 4 s-labeled nodes being the same. Edges are labeled as "capacity/lower bound".

For each customer edge e, we introduce an arc (s, v_e) and the flow on this arc, denoted by $c'(e)$, represents the utility that v perceives as coming from e. This includes the payment $c(e)$ given to v, as well as values for any active demand d originating at v, i.e., $d \in D^{end}(e)$. Similarly, if v is paying a penalty $p_v(e)$, then v has extra incentive to cut e. Thus, $c'(e) = c(e) + \sum_{d \in D^{end}(e)} (\lambda_v(d) - 1) - p_v(e)$ is the utility that v sees coming from e. Note that a given flow $c'(e)$ uniquely determines the payment $c(e)$, and vice versa.[3] Similarly for each provider edge, we add an edge (c_f, s) where the flow $c'(f)$ on this edges represents the utility v perceives as leaving on a provider edge f. For a given payment $c(f)$, we have the equality $c'(f) = c(f) - \sum_{d \in D^{end}(f)} (\lambda_v(d) - 1) - p_v(f)$.

Transit: For any active demand $d \in D(v) - D^{end}(v)$, we assign one of the edges in $P(d)$ incident to v to *own* the demand. A provider/peer edge will never own a demand, so if d uses a customer edge e as well as a provider/peer edge, then we set e to be the owner. Otherwise, it uses two customer edges e, f and we pick an owner arbitrarily. Now for each customer edge e, we include an arc (v_e, s) with a lower bound equal to the number of demands it owns. These arcs represent utility lost due to transit.

Bounded Lending: For any active pair of edges e, f incident to v, let $P(e, f)$ denote the set of penalty-enabled demands where v would have to pay a penalty to the other endpoint of e if f were not active. Then, we form an edge (v_e, v_f), with capacity $\sum_{d \in P(e,f)} (\lambda_u(d) - 1) - O(e, f)$, where u is the endpoint of d below e, and $O(e, f)$ is the number of active demands in $P(e, f)$ that e owns. Thus edges are allowed to transfer money to each other, but only after paying for the transit of the demands they own. We call this the *bounded lending constraint.*

Finally, to represent profit which is kept by node v, for every node v_e we add an extra edge (v_e, s), with lower bound 0 and infinite capacity. Figure 4 illustrates this point with arcs into the "profit copy" of node s. Note that $c'(e)$'s can actually be negative if $p_v(e)$ is large. Thus in fact, we would also include reverse arcs (v_e, s) in the construction and let $c'(e)$ be the flow on (s, v_e) minus that on (v_e, s).

Lemma 6.2 *Node v is stable with payment vector c if and only if there exists a circulation where the flow on corresponding edges is c'.*

One can hook up the circulation gadgets above to form a graph $G(S)$, where feasible flows uniquely determine Nash equilibrium payments with active edges S, and vice versa.

Theorem 6.3 *An edge set S induces a Nash equilibrium with payment vector c if and only if $G(S)$ has a feasible flow with values c' on corresponding edges.*

[3] If we end up with some $c(e) < 0$, it just means that v is stable, even with $c(e) = 0$.

Corollary 6.4 *For any active edge set S, we can polytime compute an NE payment vector c for S (or determine that none exists) that optimizes any linear objective function.*

Acknowledgments

We would like to thank Martin Pál and Jon Kleinberg for productive and illuminating discussions. We are grateful to Jennifer Rexford and Francis Zane for several helpful insights.

References

[1] E. Anshelevich, A. Dasgupta, É. Tardos, T. Wexler. Near-optimal network design with selfish agents. *STOC*, 2003.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. *FOCS*, 2004.

[3] E. Anshelevich, B. Shepherd, G. Wilfong. Strategic Network Formation through Peering and Service Agreements. See `www.cs.rpi.edu/~eanshel`

[4] V. Bala, S. Goyal. Self-organization in communication networks. McGill University Technical Report, 1996.

[5] V. Bala, S. Goyal. A noncooperative model of network formation. *Econometrica*, 68, pp. 1181–1229, 2000.

[6] F. Bloch, M.O. Jackson. The Formation of Networks with Transfers among Players. *J. Economic Theory*, to appear.

[7] R. Cáceres, N. Duffield, A. Feldmann, J.D. Friedmann, A. Greenberg, R. Greer, T. Johnson, C.R. Kalmanek, B. Krishnamurthy, D. Lavelle, P.P. Mishra, J. Rexford, K.K. Ramakrishnan, F.D. True, J.E. van der Merwe. Measurement and analysis of IP network usage and behaviour. *IEEE Communications Magazine*, pp. 144–151, May 2000.

[8] J. Corbo, T. Petermann. Selfish peering and routing in the Internet. CoRR cs.GT/0410069, 2004.

[9] J. Corbo, D. C. Parkes. The price of selfish behavior in bilateral network formation. *PODC*, pp. 99–107, 2005.

[10] J. Corbo, D. C. Parkes. The Role of Bilateral Consent in Strategic Network Formation. Harvard Working Paper Series.

[11] T. Erlebach, A. Hall, A Panconesi, D. Vukadinović. Cuts and Disjoint Paths in the Valley-Free Path Model of Internet BGP Routing. *Proceedings of CAAN,* pp. 49–62, 2004.

[12] T. Erlebach, A. Hall, T. Schank. Classifying Customer-Provider Relationships in the Internet. *Proceedings in Informatics*, pp. 52–64, September 2002.

[13] A. Fabirikant, A. Luthra, E. Maneva, C.H. Papadimitriou, S. Shenkar. On a network creation game. *PODC*, 2003.

[14] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenkar. A BGP-based Mechanism for Lowest-Cost Routing. *Distributed Computing 18*, pp. 61-72, 2005.

[15] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, F. True. Deriving traffic demands for operational IP networks: methodology and experience. *IEE/ACM Trans. on Networking* vol. 9, No. 3, pp. 265–279, June 2001.

[16] L. Gao, J. Rexford. Stable internet routing without global coordination. *Proceedings of ACM SIGMETRICS '00*, Santa Clara, CA, June 2000.

[17] A. Gomes. Multilateral contracting with externalities. *Econometrica*, 73(4), pp. 1329–1350, 2005.

[18] G. Gopalakrishnan, B. Hajek. Do greedy autonomous systems make for a sensible internet? *Conference on Stochastic Networks*, June 19-24, 2002.

[19] M. Goemans, D. Williamson. The primal-dual method for approximation algorithms and its application to network design problems. In *Approximation Algorithms for NP-Hard Problems,* ch. 4, pp 144-191. PWS Publishing Company, 1997.

[20] A. Greenberg, A. Shaikh. Operations and Management of IP Networks: What Researchers Should Know. *SIGCOMM Tutorial*, 2005.

[21] T. Griffin, G. Wilfong. On the Correctness of IBGP Configuration. *Proceedings of SIGCOMM,* pp. 17–29, 2002.

[22] P. Hall. On Representatives of Subsets. *Journal of London Mathematical Society*, 10, pp. 26–30, 1935.

[23] L. He, J. Walrand. Dynamic provisioning of service level agreements between interconnected networks. *Conference on Stochastic Networks*, June 19-24, 2002.

[24] H. Heller, S. Sarangi. Nash networks with heterogeneous agents. *Virginia Tech. Working Paper Series*. E-2001-1, 2001.

[25] B.E. Hermalin, M.L. Katz. Network interconnection with two-sided user benefits. *Working paper available:* `http://faculty.haas.berkeley.edu/hermalin/`

[26] `http://hosting.bellsouth.net/bellsouthhosting/s/s.dll?spage=cg/news/service_level.htm`

[27] `http://www.verio.com/access/pricing.cfm`

[28] `http://www.xo.com/products/smallgrowing/internet/dia/oc3.html`

[29] M. Jackson, A survey of models of network formation: stability and efficiency. *Group Formation in Economics: Networks, Clubs and Coalitions*, eds. G. Demange and M.Wooders, Cambridge Univ. Press.

[30] R. Johari, J. Tsitsiklis. Routing and Peering in a Competitive Internet. *IEEE Conference on Decision and Control*, 2004.

[31] R. Johari, S. Mannor, and J. Tsitsiklis. A contract-based model for directed network formation. To appear in *Games and Economic Behavior.*

[32] E. Koutsoupias, C. Papadimitriou. Worst-case equilibria. *Proceedings of STACS,* pp. 404–413, 1999.

[33] J.J. Laffont, J.S. Marcus, P. Rey, J. Tirole. Internet connection and the off-net pricing principle, *The RAND Journal of Economics*, vol. 34, n. 2, pp. 370–390, 2003.

[34] J. Larrosa, F. Tohmé. Network formation with heterogeneous agents. Economics Working Paper Archive EconWPA, Microeconomics series, No. 0301002, 2003.

[35] B. Shepherd, G. Wilfong. Multilateral Transport Games. *Proceedings of INOC*, pp. B2-378–B2-385, 2005.

[36] S. Shakkottai, R. Srikant. Economics of network pricing with multiple ISPs. *INFOCOM*, 2005.

[37] H. Wang, H. Xie, Y.R. Yang. On the stability of interdomain inbound traffic engineering. *SIGCOMM Poster Session*, 2005.