

# Incentive-Compatible Interdomain Routing with Linear Utilities\*

Alexander Hall<sup>1</sup>, Evdokia Nikolova<sup>2</sup>, and Christos Papadimitriou<sup>1</sup>

<sup>1</sup> UC Berkeley, USA, alex.hall@gmail.com, christos@cs.berkeley.edu

<sup>2</sup> MIT CSAIL, USA, nikolova@mit.edu

**Abstract.** We revisit the problem of incentive-compatible interdomain routing, examining the, quite realistic, special case in which the autonomous systems' (ASes') utilities are linear functions of the traffic in the incident links, and the traffic leaving each AS. We show that incentive-compatibility towards maximizing total welfare is achievable efficiently, and, in the uncapacitated case, by an algorithm that can be implemented by BGP, the standard protocol for interdomain routing.

## 1 Introduction

The Internet is in many ways a mysterious object, a complex wonder which we must approach with the same puzzled humility with which neuroscientists approach the brain and biologists the cell. Even at the most basic level of routing, for example, it is not clear at all how and why the approximately 20,000 independent, and presumably selfish, autonomous systems (ASes) cooperate to provide connectivity between any two of them. The problem is quintessentially economic. ASes are known to have confidential financial agreements on how traffic between them is to be handled and paid for, and such agreements are reflected in the ways in which each AS routes traffic. We can think of the ASes as nodes of an undirected graph, with edges signifying the existence of such an agreement between the two endpoints (equivalently, the possibility of traffic routed directly between the two). In particular, ASes communicate in terms of the border gateway protocol (BGP), a flexible protocol allowing them to implement routing decisions of arbitrary complexity, by “advertising” paths to adjacent ASes, and selecting among the paths advertised by their neighbors. Hence, the Internet is in its essence an economy, a game, an arena where agents act selfishly and are affected by everybody’s decisions; consequently, one can ask of it the questions we usually ask of such systems, for example the price of anarchy, or the possibility of incentive-compatible maximization of social welfare (questions typically studied by algorithmic mechanism design [18]); in this paper we address the latter.

Indeed, starting with Feigenbaum *et al.* [10], BGP has been studied in the past under the lens of algorithmic mechanism design, and in particular in terms of

---

\* The authors were supported through: NSF grant CCF - 0635319, a gift from Yahoo! Research, a MICRO grant, the American Foundation for Bulgaria Fellowship, and a fellowship by the Swiss National Science Foundation PA002 - 113170 / 1.

the Vickrey-Clarke-Groves (VCG) mechanism (see, e.g., [17] for an introduction to mechanism design). It was noticed [10] that social welfare can be optimized in routing, if one assumes that each AS has a fixed per packet cost, via the VCG mechanism with payments; and in fact, that this can be achieved in a way that is very “BGP-friendly,” i.e., can be implemented by minimal disruption of BGP’s ordinary operation. Furthermore, it was observed that in the real Internet VCG would result in relatively very small overpayments.

In a subsequent paper [12], the problem of more realistic BGP routing was addressed in the same spirit. Each AS was assumed to have a utility for each path to the destination (assumed in this literature to be a fixed node 0), and the goal is to maximize total utility. It was shown that the problem is too hard to solve in general even with no consideration to incentive compatibility, while a special case, in which the utility of a path only depends on the next hop, is easy to solve in an incentive-compatible way, but hard to implement on BGP. To show this latter negative result, the authors of [12] formalize what it means for an algorithm to be “BGP-friendly”: roughly speaking, a local distributed algorithm with quick convergence, small storage needs, and no rippling updates in case of small parameter changes. All said, the message of Feigenbaum, Shenker and Sami [12] was that, if one models BGP routing a little more realistically, incentive compatibility becomes problematic. This negative message was ameliorated in [11], where it was pointed out that, if one further restricts the special case of next-hop utilities so that paths are required to be of a particular kind mandated by the kinds of inter-AS agreements seen in practice, called *valley-free* in this paper, BGP-friendly incentive compatibility is restored.

There is an extensive literature on BGP (see, e.g., [9, 14, 16, 20–22]). The protocol has also been examined within other game-theoretic contexts, such as with respect to network creation games, e.g., [2, 7], cooperative game theory [19], and BGP oscillation prediction [8].

In this paper we present an elementary model of BGP routing. The key feature of our model is that *path preferences are based exclusively on per packet costs and per packet agreed-upon compensation between adjacent nodes*. In other words, we look into the utilities of each path to each AS, taken as raw data in previous literature, and postulate that they are linear functions of the traffic, depending on two factors: Objective per packet costs to each AS for each incoming or outgoing link, and agreed per packet payment, positive or negative, to the AS for this link and direction. As a result, social welfare optimization becomes a min-cost flow problem, and incentive-compatibility can always be achieved in polynomial time. If there are no capacity constraints, we show (Theorem 1) that the resulting algorithm is BGP-friendly, essentially because the BGP-friendly version of the Bellman-Ford algorithm in [10] can be extended to cover this case. When capacities are present, the algorithm becomes a more generic min-cost flow computation (Theorem 2), and, as we show by a counterexample, does not adhere to the criteria of BGP-friendliness (it may not converge fast enough), even though it is still a local, distributed algorithm with modest memory requirements and no need for global data. If, on top of this, we also require that the paths be

of the “valley-free” kind suggested by the kinds of agreements between ASes one sees in practice (that is, the kind of restriction which led to tractability in [11]), the resulting algorithm solves a rather generic linear program (Theorem 3), and so local, distributed computation appears to be impossible.

## 2 Basic Model & the VCG Mechanism

We model interdomain routing as a symmetric directed network with node set  $V = \{0, 1, \dots, n\}$  and edges  $E$ , where node 0 is the given destination, assumed unique as is common in this literature. Note that we postulate that the network is symmetric, in that if  $(i, j) \in E$  then also  $(j, i) \in E$ . There are no self-loops. Each node  $i$  has a *demand* of  $k_i$  packets it wants to send to the destination. In addition, each node  $i$  has a *per packet value*  $v_{i,e}$  (sometimes also denoted as  $v_i(e)$ ) for each of its incident edges  $e$ , and a value  $\pi_i$  for each of its packets that gets delivered. The *cost* of an edge  $e = (i, j) \in E$  is the negative of the sum of values of  $i$  and  $j$  for it,  $p_e = -(v_{i,e} + v_{j,e})$ .

We denote by  $\theta_i$  the *type* of node  $i$ , that is the collection of values for its incident edges and its value per packet delivery. Denote by  $\theta$  the vector of all node types and by  $\theta_{-i}$  the vector of all node types except that of node  $i$ .

If  $F$  is an integer-valued flow through this network, with sink 0 and sources at all other nodes with the given demands, then the utility of each node  $i \neq 0$  from this flow is  $v_i(F, \theta_i) = \sum_j v_i(i, j)F(i, j) + \sum_j v_i(j, i)F(j, i) + \pi_i F_i$ , where by  $F_i = \sum_j F(i, j) - \sum_j F(j, i)$  we denote the flow out of  $i$ , assumed to be at most  $k_i$ . The total welfare of  $F$  is  $W(F) = \sum_{i \in V \setminus \{0\}} \pi_i F_i - \sum_{e \in E} p_e F(e)$ . Let  $F^*(\theta)$  be the optimum, with respect to  $W$ , flow for types  $\theta$ ; we denote  $W(F^*(\theta))$  simply by  $W^*(\theta)$ ;  $W^*(\theta_{-i})$  is the welfare of the optimum flow when node  $i$  is deleted from the network. We assume initially that all capacities are infinite, which implies that the optimum flow is the union of  $n$  or fewer flow-weighted source-to-sink shortest paths; this assumption is removed in Section 2.2.

### 2.1 VCG Mechanism.

Notice that in order to compute the optimum flow we need to know the types of all players; the difficulty is, of course, that the type of player  $i > 0$  is known only to player  $i$ , who is not inclined to publicize it in the absence of appropriate incentives. The *VCG mechanism* for this problem incentivizes the players to reveal their true types, and thus participate in a socially optimum flow, by making payments to them. Let  $|a|_{[b \geq 0]} = a$  for  $b \geq 0$  and  $|a|_{[b \geq 0]} = 0$  otherwise. Consider in particular the following transfers for each node (negative for payments made by the node and positive for payments received by the node).

$$\begin{aligned} t_i(\theta) &= \left[ \sum_{j \neq i} v_j(F^*(\theta), \theta_j) \right] - \left[ \sum_{j \neq i} v_j(F^*(\theta_{-i}), \theta_j) \right] \\ &= \sum_{j \neq i} k_j \cdot (|\pi_j - P_j^{-i}|_{[\pi_j - P_j \geq 0]} - |\pi_j - P_{j,-i}|_{[\pi_j - P_{j,-i} \geq 0]}), \end{aligned}$$

where  $P_{j,-i}$  is the cost of the cheapest path from  $j$  to 0 which does not go through node  $i$ .  $P_j^{-i}$  is the cost of the cheapest path  $p_j$  from  $j$  to 0 without taking costs potentially incurred by  $i$  into account: if  $i \notin p_j$ ,  $P_j^{-i} = P_j$ , otherwise  $P_j^{-i} = P_j + (v_{i,e_1} + v_{i,e_2})$  with  $e_1, e_2 \in p_j$  denoting the edges incident to  $i$ .

The proof that these transfers lead to truthful reporting is the corresponding proof about the Groves mechanism in [17] specialized to the current situation. We repeat it here for completeness:

*Proof.* Suppose truth is not a dominant strategy for some node  $i$ , that is the node gets higher utility by reporting a collection of values  $\hat{\theta}_i$  different from his true values  $\theta_i$  when the other nodes report  $\theta_{-i}$ . The utility of the node is its welfare plus the transfer imposed by the mechanism:  $v_i(F^*(\hat{\theta}_i, \theta_{-i}), \theta_i) + t_i(\hat{\theta}_i, \theta_{-i}) > v_i(F^*(\theta), \theta_i) + t_i(\theta_i, \theta_{-i})$ .

Substituting the form of the transfer on both sides and canceling identical terms, we get  $v_i(F^*(\hat{\theta}_i, \theta_{-i}), \theta_i) + \left[ \sum_{j \neq i} v_j(F^*(\hat{\theta}_i, \theta_{-i}), \theta_j) \right] > v_i(F^*(\theta), \theta_i) + \left[ \sum_{j \neq i} v_j(F^*(\theta), \theta_j) \right] \Leftrightarrow W(F^*(\hat{\theta}_i, \theta_{-i}), \theta) > W(F^*(\theta), \theta)$ . The last inequality contradicts the fact that  $F^*(\theta)$  is the welfare maximizing choice of paths (*i.e.*, the least cost paths) for node types  $\theta$ .  $\square$

## 2.2 The Model with Capacities

In this subsection we consider the same basic model, with the addition that each edge  $e$  has a corresponding capacity  $c_e$ . We would like to find a min-cost (multicommodity) flow from all nodes to the sink 0, satisfying the demands of the nodes. We can transform the problem to an equivalent one by adding a new node—a supersource, which is connected to each node  $j$  via an edge of cost  $-\pi_j$  and capacity equal to the demand  $k_j$  at node  $j$ .

Call the resulting min-cost flow with known types  $\theta$  by  $F^*(\theta)$ , and denote the min-cost flow in the graph with node  $i$  removed as  $F^*(\theta_{-i})$ . We can now get a VCG mechanism similar to the one in the basic model above. As before, the total welfare is  $W(F^*(\theta), \theta) = \sum_i v_i(F^*(\theta), \theta_i)$ , where  $v_i(F^*(\theta), \theta_i)$  is the value of the flow from  $i$  (more precisely, from the supersource through  $i$ ) to 0. Similarly, the VCG mechanism is specified by the transfers

$$t_i(\theta) = \left[ \sum_{j \neq i} v_j(F^*(\theta), \theta_j) \right] - \left[ \sum_{j \neq i} v_j(F^*(\theta_{-i}), \theta_j) \right]$$

and a proof of truthfulness of the mechanism follows as before.

## 3 Economic Relationships

The economic relationships between individual ASes in the Internet severely influence the paths which can be taken in the BGP graph. So far we assumed that all paths which are present in the underlying undirected graph (there is an

edge between two ASes, if they are connected by a physical link) are valid. In reality this is not the case. Routing policies which are based on the economic relationships between ASes forbid many paths which theoretically exist. Inferring these economic relationships and investigating the resulting consequences for the connectivity of the BGP graph have attracted a large amount of scientific interest recently, see, e.g., [1, 3, 4, 6, 14, 15, 22].

Below we will give a detailed description of the valley-free path model which classifies the prohibited paths in the BGP graph.

*The Valley-Free Path Model.* In this model there are basically three different types of relationships a pair of connected ASes can be in: either *customer-provider*, in which the customer pays the provider to obtain access to the Internet, or *peer-peer*, in which both peers agree on mutually routing traffic of their customers for each other free of charge, or *sibling*, in which both siblings agree on mutually routing any traffic for each other free of charge. Note that an individual AS may take several roles—as customer, provider, sibling, or peer—simultaneously; it can for instance be a customer of one AS and at the same time a provider for another AS.

In the following for ease of exposition we will focus on customer-provider relationships only. The other types of relationships (peer-peer, sibling) can be incorporated easily, as we will note in Section 4.3.

We call a directed graph  $G = (V, E)$  a *ToR graph*, if it contains no self loops and the edge directions describe the economic relationships. If the AS  $v$  is a customer of a provider AS  $w$ , we direct the edge between  $v$  and  $w$  towards  $w$ . This follows the terminology of [22].

In practice routing is done in the following way. If AS  $w$  is a provider of  $v$  (i.e.  $(v, w) \in E$ ) it announces all its routes to  $v$ , but AS  $v$  on the other hand only announces its own routes and the routes of its customers to  $w$ . In other words, the customer  $v$  essentially advertizes only its incoming links to the provider  $w$ . The idea behind this is that  $v$  pays  $w$  for the connection and thus is not inclined to take over “work” for  $w$ . This would happen if  $v$  also announced the routes it has from other providers. Then it would potentially have to donate bandwidth to packets that arrive from the provider  $w$ , only to proceed to another provider.

This leads to the model proposed in [22] that a path  $p$  is *valid* if and only if it consists of a sequence of customer-provider edges ( $\bullet \rightarrow \bullet$ ) followed by a sequence of provider-customer edges ( $\bullet \leftarrow \bullet$ ). The first part, containing only customer-provider edges, is also called the *forward part* of  $p$ . The last part, containing only provider-customer edges, is called the *backward part* of  $p$ . It is easy to see that the following is an equivalent definition of the validity of a path:

A path  $p = v_1, e_1, v_2, e_2, \dots, e_{r-1}, v_r$  in the ToR graph  $G$  is a *valid*  $v_1$ - $v_r$ -*path* in  $G$ , if and only if there is no inner node  $v_i$  of  $p$  for which  $e_{i-1}$  and  $e_i$  are outgoing edges of  $v_i$ .

If such an inner node—one which does have this property—exists, it is called a *valley*. The intuition behind this name is that outgoing edges point “upwards”, out of the valley. In the literature the situation that a path contains a valley is

also called an *anomaly*. A flow which only uses valley-free paths we call a *valid* or *valley-free* flow.

*The VCG mechanism.* The transfers can be specified as in Section 2.2 for the model with capacities. The only difference is that all flows (i.e.,  $F^*$  and  $F_{-i}^*$  for all  $i$ ) must be valley-free (and may be fractional).

## 4 Distributed Computation of VCG Payments

It is of great interest to determine to what extent the payments  $t_i(\theta)$  can be computed not only efficiently, but in a distributed manner which is “BGP-friendly,” that is, compatible with current usage of the BGP protocol. In [12] this concept of “BGP-friendliness” was formalized as three requirements:

1. The algorithm should converge in a number of rounds that is proportional to the diameter of the graph, and not its size.
2. Only local data should be needed.
3. No rippling updates should be needed as data changes.

Here we relax requirement (1) to a number of rounds that is proportional to the diameter times  $R$ , where  $R$  is the ratio between the largest and smallest edge cost. This is necessary (also in [12], where the stricter version is used by oversight) because the computed shortest path, whose length is the upper bound on convergence time, may be longer than the diameter. We do not bother to formalize here the second and third requirement (the reader is referred to [12]) because our algorithms either trivially satisfy any conceivable version, or fail to satisfy (1). As it turns out, this important aspect of BGP-friendliness sets the basic model apart from the model with capacities. In both cases the implementation is quite simple and makes only modest use of local resources. But only in the former case the strict conditions on the convergence time are fulfilled.

### 4.1 Basic Model

For the basic model it is easy to adapt the approach presented by Feigenbaum et al. [10]. BGP is a *path-vector* protocol which computes the lowest-cost paths (LCPs) in a sequence of stages. In a stage each node in the network sends all the LCPs it knows of to its neighbors. It also receives LCPs from its neighbors. If these contain shorter paths than the ones it has currently stored, it updates the list of its own LCPs. This basically corresponds to a distributed computation of all shortest paths via the Bellman-Ford algorithm. The computation terminates after  $d$  stages and involves  $O(nd)$  communication on any edge, where  $d$  denotes the maximum number of edges on an LCP.

Feigenbaum et al. give an interesting and easy to implement extension of the path-vector protocol which computes not only the lowest cost paths, but at the same time the lowest cost paths which do not traverse a given node  $i$ . These two quantities are then used to compute the payments for node  $i$ . This increases

the number of stages and communication needed to  $d'$  and  $O(nd')$ , respectively. Here  $d'$  denotes the maximum number of edges on an LCP avoiding node  $i$ , over all nodes  $i$  for which  $G \setminus \{i\}$  is still connected. Feigenbaum et al. argue that this is still an acceptable convergence time.

The only difference of the approach in [10] to ours is that the per packet values in our model are given individually for each edge and node, i.e., as  $v_{i,e}$ , and not only as one total value per node.<sup>3</sup> Hence, it is easy to adapt their method to compute the values  $P_j$  and  $P_{j,-i}$ , for  $j, i \in \{1, \dots, n\}$ , which is all we need to compute  $t_i(\theta) = \sum_{j \neq i} (P_{j,-i} - P_j^{-i})k_j$ . Note that the partial path cost  $P_j^{-i}$  can be easily derived from the cost  $P_j$  of the cheapest path from  $j$  to the sink.

Let  $\text{diam}'(G)$  denote the maximum diameter of  $G \setminus \{i\}$  (as  $d'$ , measured in number of edges) over all nodes  $i$  for which  $G \setminus \{i\}$  is still connected. Since  $d' \leq \text{diam}'(G) \cdot R$ , where  $R$  is the ratio between the largest and smallest edge cost, we obtain the following theorem.

**Theorem 1.** *In the basic per packet utility model without capacity constraints (described in Section 2) the Vickrey-Clarke-Groves allocation and payments can be computed in a distributed, BGP-friendly manner. The computation converges in  $O(\text{diam}'(G) \cdot R)$  rounds of communication.*

## 4.2 Model with Capacities


Instead of lowest cost paths and lowest cost paths avoiding node  $i$ , we now need to know a min-cost flow  $F(\theta)$  and a min-cost flow  $F_{-i}(\theta)$  avoiding node  $i$  for each of the payments  $t_i(\theta)$ ,  $i \in \{1, \dots, n\}$ . In the following we will explain how to compute  $F(\theta)$  in a distributed fashion. The flow  $F_{-i}(\theta)$  can be computed correspondingly by blocking node  $i$ . Therefore, altogether  $(n+1)$  flow computations are performed, one for  $F(\theta)$  and  $n$  for the  $F_{-i}(\theta)$ ,  $i \in V \setminus \{0\}$ .

We assume the sink 0 controls all the computations: it chooses which node is blocked (in the  $F_{-i}(\theta)$  case), it selects paths to send flow along together with the corresponding amounts, and it recognizes when a min-cost flow computation is finished. These all are computationally simple tasks. The only intensive computations needed will be those to obtain the shortest paths with respect to certain costs and where certain edges may be blocked. These will be done in a distributed manner applying the standard distributed Bellman-Ford algorithm, which is used by BGP as mentioned above.

*Distributed Computation of  $\mathbf{F}(\theta)$ .* We start with a description of a simple Ford-Fulkerson approach [13] of computing a min-cost flow from the supersource to the sink via augmenting shortest paths. Then we explain how to modify it to use the Edmonds-Karp scaling technique [5].

A virtual residual graph is overlaid over the given network. The residual edge capacities and costs are derived from the original graph. The residual capacities depend on the flow present on the corresponding residual edges and thus

<sup>3</sup> This allows for more fine-granular and thus more realistic modeling.



**Fig. 1.** All edges have capacity 1, except the top edge with capacity  $l + 1$ . The edge costs are all 0, except the rightmost edge with cost 1. All nodes have a demand of 1 to be sent to node 0.

may change during the computation of the flow. Each node keeps track of flow values on residual edges incident to it.

Consider an original pair of directed edges  $(i, j)$  and  $(j, i)$  with costs  $p_{(i,j)}$  and  $p_{(j,i)}$ . We assume the costs to be greater or equal to 0. Let  $f_{(i,j)}$  and  $f_{(j,i)}$  denote the flow amounts on these edges, only one of which may be greater 0. Otherwise, a circular flow of  $\min(f_{(i,j)}, f_{(j,i)})$  is subtracted from both without increasing the costs. The residual capacities are set to  $c'_{(i,j)} = c_{(i,j)} - f_{(i,j)}$  and  $c'_{(j,i)} = c_{(j,i)} - f_{(j,i)}$ . Additionally, we add the virtual edges  $\overleftarrow{(i, j)}$  and  $\overrightarrow{(j, i)}$  with capacities  $c_{\overleftarrow{(i, j)}} = f_{(j,i)}$  and  $c_{\overrightarrow{(j, i)}} = f_{(i,j)}$  and costs  $p_{\overleftarrow{(i, j)}} = -p_{(i,j)}$  and  $p_{\overrightarrow{(j, i)}} = -p_{(j,i)}$ . Flow sent onto these edges is subtracted from the corresponding flow on the edge in the opposite direction. Finally, for each  $i \in V \setminus \{0\}$  a virtual edge is added from the supersource to node  $i$  with cost  $-\pi_i$ .

The algorithm now proceeds as follows, steps 2-4 comprise a phase.

1. For each node  $i \in V$  initialize the flow values  $f_{(i,j)} = f_{(j,i)} = 0$  of all incident residual edges. Update the local capacities as described above.
2. Compute the shortest paths in the current residual graph only considering edges with capacities greater than 0. Do this with the distributed Bellman-Ford algorithm, adapting the BGP implementation. Modify the algorithm to also forward the bottleneck capacity of each path.
3. The sink checks the min-cost path to the supersource. If the cost is  $\geq 0$ , we are done. Otherwise send a flow corresponding to the bottleneck capacity along the path. This is done by sending a message along the path, which notifies the contained nodes to update their local flow values (and thus capacities).
4. Continue at step 2 with the updated residual graph.

*Time to Converge, Improvements.* Each phase consists of a (re)computation of the shortest paths in Step 2. Unfortunately, in the capacitated case the rounds of communication for a shortest paths computation is not bounded by  $d$  or  $d'$  anymore. It may actually take up to  $n$  rounds of communication, as the example at the end of this subsection shows.

Let  $C = \max\{c_e | e \in E\}$  be the maximum capacity. The algorithm finishes in  $O(|E| \cdot C)$  phases. This can be improved to  $O(n \cdot \log C)$  by applying the



following well-known scaling technique. A variable  $\Delta$  is introduced and initialized to  $2^{\lceil \log C \rceil - 1}$  in step 1. In step 2 only edges with capacity  $\geq \Delta$  are considered. In step 3,  $\Delta$  is updated to  $\Delta/2$ , if no more negative cost paths are found (unless  $\Delta = 1$ , then we are done). The updated  $\Delta$  is broadcast to all nodes.

As mentioned, with  $(n + 1)$  such flow computations we can compute all node payments  $t_i(\theta)$ . Altogether this yields the following theorem.

**Theorem 2.** *In the per packet utility model with capacity constraints (described in Section 2.2) the VCG allocation and payments can be computed in a distributed manner. The computation converges in  $O(n^3 \cdot \log C)$  rounds of communication.*


*Shortest Paths Computation.* Unfortunately, the number of rounds of communication to compute the shortest path cannot be bound by  $d$  (or  $d'$ ) anymore. Figure 1 shows an example where the shortest path in the residual graph has length  $n - 2$ , whereas the number of hops in the corresponding LCP in the original graph is 2. Assume that all nodes have already (virtually) sent their flow through the residual graph except node 1 which is selected last. Since the nodes are indistinguishable, we may assume this. The only path remaining in the residual graph is the one at the bottom of length  $n - 2$ , since the capacities of all other edges (except  $(1, 2)$ ) are fully saturated by flow sent to the sink via node 2. This compares to the LCP from node 1 over node 2 directly to node 0 with only two edges.

### 4.3 Model with Economic Relationships

In the following we will explain the two-layer graph, a helpful notion which was originally suggested in [6]. With the help of the two-layer model it will be easy to see that one can compute min-cost valley-free flows as needed in our model with capacities introduced in Section 2.2.

*The Two-Layer Model.* From a ToR graph  $G = (V, E)$  and source, sink  $s, t \in V$  we construct a *two-layer model*  $H$ , which is a directed graph, in the following way (see Figure 2 for an example). Two copies of the graph  $G$  are made, called the *lower* and the *upper layer*. In the upper layer all edge-directions are reversed. Every node  $v$  in the lower layer is connected with an edge to the corresponding copy of  $v$ , denoted  $v'$ , in the upper layer. The edge is directed from  $v$  to  $v'$ . Finally, we obtain the two-layer model  $H$  by identifying the two  $s$ -nodes (of lower and upper layer) and also the two  $t$ -nodes, and by removing the incoming edges of  $s$  and the outgoing edges of  $t$ .

A valid path  $p = v_1 \cdots v_r$  in  $G$  with  $v_1 = s$  and  $v_r = t$  is equivalent to a directed path in  $H$  in the following way. The forward part of  $p$ , that is the part containing all edges  $(v_i, v_{i+1}) \in p$ , is routed in the lower layer. Then there is a possible switch to the upper layer with a  $(v, v')$ -type edge (there can be at most one such switch for each path). The backward part of  $p$  is routed in the upper layer. In other words for each original edge  $(v_{i+1}, v_i) \in p$  the corresponding edge  $(v'_i, v'_{i+1})$  of the upper layer is traversed. If there is only a forward (respectively



**Fig. 2.** A path in the ToR graph  $G$  and the corresponding path in the two-layer model  $H$ . ( $G'$  is  $G$ , excluding  $s$  and  $t$ .)

backward) part of  $p$ , then the corresponding path in  $H$  is only in the lower (respectively upper) layer.

This definition of the two-layer model can easily be extended to the case of multiple sources. Note that a peer-peer relationship between two nodes  $v, u \in V$  can be incorporated by adding the edges  $(v, u')$  and  $(u, v')$  from lower to upper layer (reflecting that at most one peer-peer edge is allowed between the forward and the backward part of a path). Similarly, a sibling relationship between two nodes  $v, u \in V$  can be incorporated by adding the symmetric edges  $(v, u)$ ,  $(u, v)$ ,  $(v', u')$ , and  $(u', v')$  in both layers (reflecting that sibling edges are allowed at arbitrary points in a path).

*Min-Cost Valley-Free Flows.* By simply computing a min-cost flow in the two-layer graph it is easy to derive a valley-free flow which will have at most the cost of an optimum min-cost valley-free flow. The edge capacities may be violated by at most a factor of two though, since each edge may be used twice: once in the upper and once in the lower layer. Note that such a min-cost flow could be computed in a distributed fashion by slightly modifying the approach described in Section 4.2.

This approximate solution cannot be used to compute the VCG allocation and payments though. To this end, we need the optimal solution. The latter can be computed with the help of a standard LP flow formulation with added constraints to bound the joint flow on the upper and lower layer edges. In other words, for each edge  $(v, u) \in E$  in the original ToR graph, we add a joint capacity constraint for  $(v, u)$  and  $(u', v')$  in the two-layer model.

**Theorem 3.** *In the per packet utility model with capacity constraints and economic relationships (described in Section 3) the VCG allocation and payments can be computed in polynomial time with an LP based approach.*

Note that the existence of an optimal algorithm based on augmenting paths seems unlikely. Usually, for integral capacities such algorithms aim at computing an optimal integral solution, i.e., for unit capacities a solution would consist of edge disjoint paths. However, computing the maximum number of disjoint valley-free paths between two nodes  $s, t$  is inapproximable within a factor of  $(2 - \epsilon)$ , unless  $P = NP$  [6].

## 5 Conclusions and Open Problems

Despite the fact that incentive compatibility for BGP routing had been known to be problematic in general, as well as for several apparently realistic special cases, we have identified one important special case of practical importance, namely the one in which path utilities depend on local per packet costs as well as delivery values. In this case incentive compatibility is achievable through payments which can be computed efficiently and in a BGP-compatible way; adding capacities and the “valley-free” constraint for paths makes incentives harder to compute in a BGP-compatible way, but still tractable.

Regarding the latter point, in this work we have simply pointed out that the algorithms we devised for VCG incentive computation are not implementable in a BGP-compatible way; it would be interesting to actually prove that this is inherent to the problem, i.e., to prove a lower bound on the convergence time of any algorithm for the min-cost flow problem and its valley-free constrained case.

Our model for path utilities is suggestive of a more general project for understanding BGP routing: We postulate that each directed edge in and out of every node has a value for this node, depending on the cost to this node, as well as agreed upon payments to or from its neighbors, for each packet sent or received along this edge. Suppose that the graph, as well as the demand, and per packet cost and delivery value of each node, are given. A game is thus defined in which strategies are payment agreements between neighbors, and the utility to each node is the one obtained by our model of BGP min-cost routing. This game is thus a very realistic network creation game, with special emphasis on BGP routing. The quality of equilibria compared to the social optimum (i.e., the price of anarchy and its variants) for this game would be a most interesting research direction. The social optimum is, of course, the min-cost flow with only costs and delivery values taken into account. Further, such a model would allow one to study how inter-AS agreements can depend on the underlying fundamentals of each AS, such as costs, delivery value, demand, and position in the network.

## References

1. D. Achlioptas, A. Clauset, D. Kempe, and C. Moore. On the bias of traceroute sampling; or, power-law degree distributions in regular graphs. In *STOC'05*, pages 694–703, 2005.
2. E. Anshelevich, B. Shepherd, and G. Wilfong. Strategic network formation through peering and service agreements. In *FOCS'06*, pages 77–86, Washington, DC, USA, 2006. IEEE Computer Society.

3. P. Barford, A. Bestavros, J. Byers, and M. Crovella. On the marginal utility of deploying measurement infrastructure. In *ACM SIGCOMM Internet Measurement Workshop*, November 2001.
4. G. Di Battista, T. Erlebach, A. Hall, M. Patrignani, M. Pizzonia, and T. Schank. Computing the types of the relationships between autonomous systems. *IEEE/ACM Transactions on Networking*, 15(2):267–280, April 2007.
5. J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19(2):248–264, 1972.
6. T. Erlebach, A. Hall, A. Panconesi, and D. Vukadinovic. Cuts and disjoint paths in the valley-free path model of Internet BGP routing. In *CAAN*, pages 49–62, 2004.
7. A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. In *PODC'03*, pages 347–351, New York, NY, USA, 2003. ACM Press.
8. A. Fabrikant and C. H. Papadimitriou. The search for equilibria: Sink equilibria, unit recall games, and BGP oscillations. In *submitted manuscript*, 2007.
9. N. Feamster, J. Winick, and J. Rexford. A model of BGP routing for network engineering. In *SIGMETRICS '04/Performance '04*, pages 331–342, New York, NY, USA, 2004. ACM Press.
10. J. Feigenbaum, C. H. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. In *PODC'02*, pages 173–182, 2002.
11. J. Feigenbaum, V. Ramachandran, and M. Schapira. Incentive-compatible inter-domain routing. In *EC'06*, pages 130–139, New York, NY, USA, 2006. ACM Press.
12. J. Feigenbaum, R. Sami, and S. Shenker. Mechanism design for policy routing. *Distrib. Comput.*, 18(4):293–305, 2006.
13. L. R. Ford and D. R. Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6:419–433, 1958.
14. L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Trans. Networking*, 9(6):733–745, Dec 2001.
15. R. Govindan and A. Reddy. An analysis of Internet inter-domain topology and route stability. In *IEEE INFOCOM'97*, pages 850–857, April 1997.
16. T. G. Griffin and G. Wilfong. An analysis of BGP convergence properties. In *SIGCOMM'99*, pages 277–288, New York, NY, USA, 1999. ACM Press.
17. A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
18. N. Nisan and A. Ronen. Algorithmic mechanism design. In *STOC'99*, pages 129–140, New York, NY, USA, 1999. ACM Press.
19. C. H. Papadimitriou. Algorithms, games, and the Internet. In *STOC'01*, pages 749–753, 2001.
20. Y. Rekhter and T. Li. A border gateway protocol 4 (bgp-4), 1995.
21. J. W. Stewart. *BGP4: Inter-Domain Routing in the Internet*. Addison-Wesley, 1998.
22. L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *IEEE INFOCOM'02*, 2002.