

Internet Economics: The use of Shapley value for ISP settlement*

Richard T.B. Ma
Columbia University
tbma@ee.columbia.edu

Dah-ming Chiu
The Chinese University of HK
dmchiu@ie.cuhk.edu.hk

John C.S. Lui
The Chinese University of HK
cslui.cse.cuhk.edu.hk

Vishal Misra
Columbia University
misra@cs.columbia.edu

Dan Rubenstein
Columbia University
danr@ee.columbia.edu

ABSTRACT

Within the current Internet, autonomous ISPs implement bilateral agreements, with each ISP establishing agreements that suit its own local objective to maximize its profit. Peering agreements based on local views and bilateral settlements, while expedient, encourage selfish routing strategies and discriminatory interconnections. From a more global perspective, such settlements reduce aggregate profits, limit the stability of routes, and discourage potentially useful peering/connectivity arrangements, thereby unnecessarily balkanizing the Internet. We show that if the distribution of profits is enforced at a global level, then there exist profit-sharing mechanisms derived from the coalition games concept of *Shapley value* and its extensions that will encourage these selfish ISPs who seek to maximize their own profits to converge to a Nash equilibrium. We show that these profit sharing schemes exhibit several fairness properties that support the argument that this distribution of profits is desirable. In addition, at the Nash equilibrium point, the routing and connecting/peering strategies maximize aggregate network profits, encourage ISP connectivity so as to limit balkanization.

1. INTRODUCTION

The Internet is composed of thousands of connected *autonomous systems* (ASes). Before transitioning to the private sector, these ASes' primary focus was to improve connectiv-

*This research was supported in part by the National Science Foundation grants CNS-0627590, ITR-0325495 and CNS-0238299. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CoNEXT'07, December 10-13, 2007, New York, NY, U.S.A.
Copyright 2007 ACM 978-1-59593-770-4/07/0012 ...\$5.00.

ity and network performance - who got paid was not the primary concern. However, in its current form, ISPs, each composed of one or more ASes, has a primary interest to maximize its own profit. Connectivity is currently implemented via bilateral agreements that are generally either a peering relationship where ISPs offer to carry one another's traffic, or customer-provider relationship where one ISP pays the other for transit [12].

These local, bilateral agreements may look beneficial from a local perspective, but from a more global perspective, they are very unappealing. ISPs will often resort to *selfish routing* such as using the hot-potato algorithm [21]. Furthermore, ISPs will often refrain from connecting to another ISP when such a connection does not increase its own profit, regardless of the benefit that the connection might provide the global system. This selfish behavior can lead to a balkanization of the Internet, with the global infrastructure dismantling into a set of networks that have varying degrees of accessibility and reachability, limiting their usefulness [8]. This balkanization inhibits the Internet's evolution toward the FCC's notion of a *universal core connecting service* [2] that can implement the mandatory functions imposed by the FCC on all telephony providers. To summarize, the lack of a more global view on the design of monetary incentives for ISPs to peer and route is limiting competition, thereby limiting technical innovation.

In this paper, we explore how to design a profit-sharing mechanism that would lead to a better engineered Internet. In other words, rather than allow ISPs to set their prices and obtain profits locally, the profit-sharing mechanism should take the collection of revenue generated by the entire network and divide this revenue "fairly" among the participating ISPs. The mechanism we implement is based on the *Shapley value* [25, 20]. This mechanism is desirable from both the global level as well as the local, ISP level. From a global perspective, the same traffic demands can be supported while increasing the aggregate network profit, and balkanization will reduce as this novel mechanism will provide more encouragement for connections. From the local perspective, the Shapley value exhibits several fairness prop-

erties that formally indicate that an ISP’s profit is proportional to its contribution to the value of the network. More specifically, our contributions are the following:

- We propose (Section 2) a novel multilateral settlement model, where customers pay for end-to-end services and ISPs collectively share the revenue for providing these services.
- We implement this settlement via a mechanism based on the Shapley value, and show that the following are achieved (Section 3):
 - **Efficiency:** The aggregate revenue delivered to the ISPs equals the aggregate payments accumulated from the customers (i.e., all funds are accounted for).
 - **Fairness:** ISPs who make greater contributions to the profit of the aggregate network receive a greater share of this profit. This general statement is specified more formally and precisely as a set of four specific properties (symmetry, balanced contribution, dummy, and strong monotonicity).
 - **Optimal Routing:** Given any fixed interconnection topology, by allowing each ISP to select routes that maximize its individual profit, the global routing topology converges to a Nash equilibrium where the aggregate profit of the system is also maximized (Section 4).
 - **Interconnection Incentives:** Each ISP’s selfish objective will encourage it to connect to other ISPs when such a connection increases the overall profit of the network, evolving the network to a better-connected, more efficient state (Section 5).
- We illustrate examples of profit distribution among ISPs including the real AS topology of Columbia University. We also simulate and compare the profits under hot-potato routing and the optimal routing resulted from our new mechanism. We show that by using the new mechanism, every ISP receives greater profit.

Our proposed mechanism is a considerable and likely controversial shift from the current bilateral settlements where an ISP’s profit are computed solely from its local interactions. This all-local property gives the ISP a false sense of independence since these profits are in fact affected by other ISPs’ decisions throughout the network. Nonetheless, this sense of independence and profit based on local perception is appealing. The two principle motivations we present for our mechanism are: (i) Our profit redistribution is *not a zero sum game*, and in fact all participating players stand to gain, and, (ii) From a social (and thus policy) perspective, our mechanism *encourages interconnection*, thereby leading to a better connected and more robust Internet.

We start off with preliminaries and a model description needed for our framework in the next section.

2. COOPERATIVE FRAMEWORK

2.1 Three Layers of the Current Internet and a Novel Two-Stage Settlement Model

We view the current Internet as three layers of bilateral interactions between ISPs as illustrated in Figure 1. At the bottom layer, pairs of ISPs decide whether or not connect. They decide the venue and type of the connections. For example, peering connections assume a symmetric traffic pattern going through the links; while customer-provider links assume asymmetric traffic flows that the provider ISP helps for transit. At the middle layer, each ISP advertises BGP routes to neighboring ISPs and decides how to route traffic efficiently to reduce its own cost. For example, hot-potato routings are often used to choose the closest egress point based on the intra-domain cost. At the top layer, end users pay their local ISPs for the services and customer ISPs pay their provider ISPs by bilateral agreements. Although all layers depend on each other, due to the limitations of bilateral interactions and selfish decisions of the ISPs, the behavior of the network from a global perspective can be highly inefficient and to a large extent unregulated.

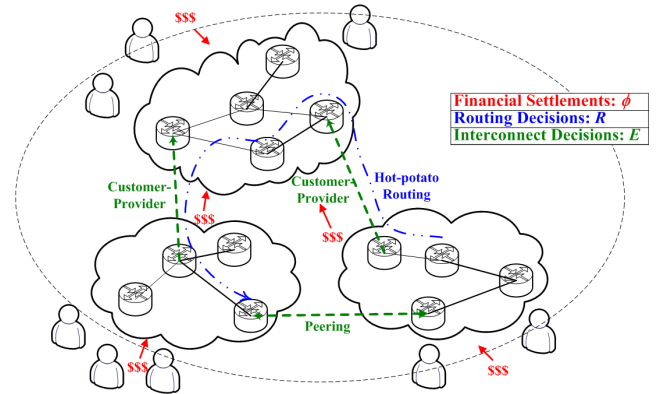


Figure 1: A view of ISP interactions of the Internet.

Unlike the existing bilateral agreements, we consider a collection of ISPs, providing end-to-end services to all their customers, as a whole. We propose a two-stage multilateral financial settlement, illustrated in Figure 2, as the following.

- S1** Customers make service agreements (charge and requirements) at their local ISPs; however, from customers’ view, the agreement is on the end-to-end service rather than a connecting and forwarding service.
- S2a** All payments from customers are collected by a multilateral profit distribution mechanism ϕ which decides the proportion of revenue each ISP receives.
- S2b** Knowing the rule of the profit distribution mechanism, each ISP makes local decisions on interconnection E and route R to maximize its profit.

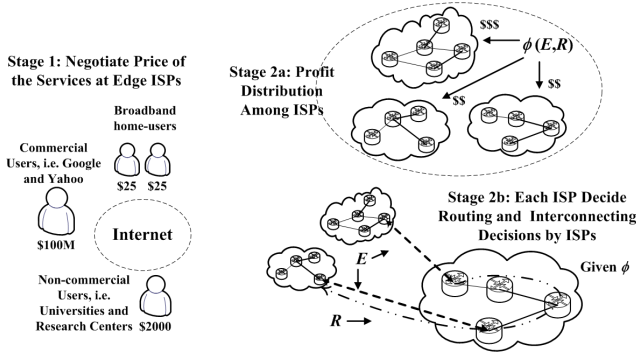


Figure 2: A two-stage multilateral settlement model.

In the first stage, the service negotiation at the edge of the network avoids the complication of buying resources from multiple ISPs along the communication paths for customers. Our pricing model is extremely general and service agreements are not restricted to service bundling, service differentiation or the pricing structure of the services. For example, services can be charged at its origin or destination. Commercial content providers might be charged more than non-profit organizations. Either usage-based pricing or flat-rate pricing can be applied. In the second stage, the mechanism $\phi(E, R)$ distributes revenues among ISPs for every possible interconnection topology E and routing decision R . Our objective is to design the profit distribution mechanism $\phi(E, R)$ that encourages selfish ISPs to interconnect extensively and route efficiently.

2.2 Network Model

We consider a network system comprised of a set of ASes. We denote \mathcal{N} as the set of ASes. $N = |\mathcal{N}|$ denotes the number of ASes in the network. We use AS and ISP interchangeably, assuming each ISP has one AS. ISPs with multiple ASes can be considered as one super-AS in our model, where any other AS connects to this super-AS if it connects to one of the ASes of the ISP. In this sense, our model does not require that the ASes of an ISP are connected.

To fairly distribute profits among ASes, we want to measure the contribution of each AS for generating those profits. In particular, we measure the profits that can be generated by subsets of the ASes. We call any nonempty subset $\mathcal{S} \subseteq \mathcal{N}$ a *coalition* of the ASes. Each coalition can be thought of as a sub-network that might be able to provide partial services to their customers. We denote v as the *worth function*, which measures the value produced by the sub-networks formed by all coalitions. In other words, for any coalition \mathcal{S} , $v(\mathcal{S})$ defines the profit generated by the sub-network formed by the set of ASes \mathcal{S} . Through the worth function v , we can measure the contribution of an AS to a group of ASes as the following.

Definition 1. The marginal contribution of AS i to a coalition $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$ is defined as $\Delta_i(v, \mathcal{S}) = v(\mathcal{S} \cup \{i\}) - v(\mathcal{S})$.

In the next section, we will describe the mechanism that uses this worth function to distribute profits among ASes. In the rest of this section, we develop the worth function for three levels of views of the network model: the AS level, the router level, and the AS peering level.

The AS level model:

We denote the AS level network system as $(\mathcal{N}, v, E_{\mathcal{N}}, R_{\mathcal{N}})$. $E_{\mathcal{N}}$ denotes the set of directed peering links between the ASes. The graph $G_{\mathcal{N}} = (\mathcal{N}, E_{\mathcal{N}})$ defines the AS level topology of the network. We denote $G_{\mathcal{S}}$ as the subgraph of $G_{\mathcal{N}}$ induced by \mathcal{S} , defined by $G_{\mathcal{S}} = (\mathcal{S}, E_{\mathcal{S}})$, where $E_{\mathcal{S}} = \{(i, j) \in E : i, j \in \mathcal{S}\}$. $G_{\mathcal{S}}$ is the AS level topology formed by the coalition \mathcal{S} . We say that AS i is *connected* to AS j by $G_{\mathcal{S}}$ if there is some $k \geq 1$ and a sequence (n_0, n_1, \dots, n_k) such that $n_0 = i, n_k = j$, and $(n_{i-1}, n_i) \in E_{\mathcal{S}}$ for $1 \leq i \leq k$.

We consider that the network provides the end-to-end data delivery services between any pair of the ASes. We denote Λ_{ij} as the required end-to-end traffic intensity from AS i to AS j . Λ_{ij} can be regarded as the volume of traffic contracted over a certain period of time for delivery. In the current Internet, the time scale for these contracts is often monthly.

We denote $R_{\mathcal{N}}$ as a *feasible route* at the AS level that achieves the traffic intensity $\{\Lambda_{ij} : i, j \in \mathcal{N}\}$ performed by all coalitions $\mathcal{S} \subseteq \mathcal{N}$. Each $R_{\mathcal{N}}(\mathcal{S})$ denotes the route on the subgraph $G_{\mathcal{S}}$ when the coalition \mathcal{S} provides the end-to-end services. Given the required traffic intensity $\{\Lambda_{ij} : i, j \in \mathcal{N}\}$, for any coalition $\mathcal{S} \subseteq \mathcal{N}$, a *feasible route* $R_{\mathcal{N}}(\mathcal{S}) = \{r_{ij}^{\mathcal{N}}(\mathcal{S}) : i, j \in \mathcal{N}\}$ defines the traffic intensity achieved on the links of $G_{\mathcal{S}}$ for every source-destination pair of ASes i and j . We define $r_{ij}^{\mathcal{N}}(\mathcal{S}, (k, l)) = 0$ if i and j are not connected by the subgraph $G_{\mathcal{S}}$ or link $(k, l) \notin E_{\mathcal{S}}$. This is because AS i is not connected to j by the coalition \mathcal{S} or the link (k, l) does not belong to the sub-network represented by $G_{\mathcal{S}}$. Otherwise, a feasible route provides the end-to-end service from AS i to AS j . We denote X_l^{ij} as the traffic intensity going through AS l . A feasible route has to satisfy the following flow conservation constraints:

$$X_i^{ij} = X_j^{ij} = \sum_{k \in \mathcal{N}} r_{ij}^{\mathcal{N}}(\mathcal{S}, (i, k)) = \sum_{k \in \mathcal{N}} r_{ij}^{\mathcal{N}}(\mathcal{S}, (k, j)) = \Lambda_{ij}.$$

$$X_l^{ij} = \sum_{k \in \mathcal{N}} r_{ij}^{\mathcal{N}}(\mathcal{S}, (k, l)) = \sum_{k \in \mathcal{N}} r_{ij}^{\mathcal{N}}(\mathcal{S}, (l, k)) \quad \forall l \in \mathcal{S} \setminus \{i, j\}.$$

Given a feasible route $R_{\mathcal{N}}$, we can also measure the aggregate traffic intensity going through each AS l as

$$X_l(\mathcal{S}, R_{\mathcal{N}}) = \sum_{i, j \in \mathcal{N}} X_l^{ij}(\mathcal{S}, R_{\mathcal{N}}). \quad (1)$$

The router level model:

In this model, we assume that the network system's router level topology is available to all ASes. We denote the router level network system as $(\mathcal{N}, v, M, E, R)$. M denote the set of routers in the network. E denotes the set of directed links connecting the routers. The graph $G = (M, E)$ defines the router level topology of the network. We denote the subset

of routers possessed by AS i as $m_i \subseteq M$. Mathematically, $\{m_i : i \in \mathcal{N}\}$ defines a partition of M , i.e. $\bigcup_{i=1}^N m_i = M$ and $m_i \cap m_j = \emptyset$ for $i \neq j$. At the router level, the topology G_S , formed by the coalition \mathcal{S} , is defined by $G_S = (M_S, E_{M_S})$, where $M_S = \bigcup_{i \in \mathcal{S}} m_i$ and $E_{M_S} = \{(i, j) \in E : i, j \in M_S\}$.

Similar to the AS level route $R_{\mathcal{N}}$, we denote R as a feasible route at the router level. Instead of having AS level traffic intensity $\{\Lambda_{ij} : i, j \in \mathcal{N}\}$, we assume that the router-to-router traffic intensity $\{\lambda_{ij} : i, j \in M\}$ is known. For any coalition $\mathcal{S} \subseteq \mathcal{N}$, a feasible route $R(\mathcal{S}) = \{r_{ij}(\mathcal{S}) : i, j \in M\}$ defines the traffic intensity achieved on the links of G_S for every source-destination pair of routers i and j . Again, we define $r_{ij}(\mathcal{S}, (k, l)) = 0$ if i and j are not connected by the subgraph G_S or link $(k, l) \notin E_{M_S}$. Otherwise, a feasible route has to satisfy the router level flow conservation constraints:

$$\sum_{k \in M} r_{ij}(\mathcal{S}, (i, k)) = \sum_{k \in M} r_{ij}(\mathcal{S}, (k, j)) = \lambda_{ij}.$$

$$\sum_{k \in M} r_{ij}(\mathcal{S}, (k, l)) = \sum_{k \in M} r_{ij}(\mathcal{S}, (l, k)) \quad \forall l \in M \setminus \{i, j\}.$$

We define \mathcal{R} as the space of all feasible routes R .




Figure 3: Router level model versus AS level model.

Both AS level and router level models can be used to describe the same network system. However, the router level model assumes that we know detailed information of the network. Given the router level information of a network, we can derive the AS level model for that network. With the router level topology (M, E) , we can construct the corresponding AS level topology $(\mathcal{N}, E_{\mathcal{N}})$ where $E_{\mathcal{N}}$ is defined by $E_{\mathcal{N}} = \{(i, j) | (k, l) \in E, k \in m_i, l \in m_j\}$. The traffic intensity at the AS level is just the aggregate router level traffic intensity, i.e. $\Lambda_{ij} = \sum_{k \in m_i, l \in m_j} \lambda_{kl} \quad \forall i, j \in \mathcal{N}$. Figure 3 illustrates an example of a router level topology and the corresponding AS level topology.

The AS peering level model:

In reality, router level information might not be available; however, AS level information might be too rough to describe the network system. For example, although the traffic patterns between ASes can often be measured and estimated as the AS level traffic intensity $\{\Lambda_{ij}\}$, the router level traf-

fic intensity $\{\lambda_{ij}\}$ could be difficult to obtain. Because each AS might consist hundreds of routers and cover a large geographical area, the AS level information, on the other hand, does not distinguish the multiple peering points between two ASes. In the AS peering level model, each AS does not need to reveal its internal router level topology. It only needs to reveal the set of edge routers (that, for instance, BGP exposes). Presumably, all the routers of an AS are connected, therefore, the set of edge router forms a logical fully-connected graph.

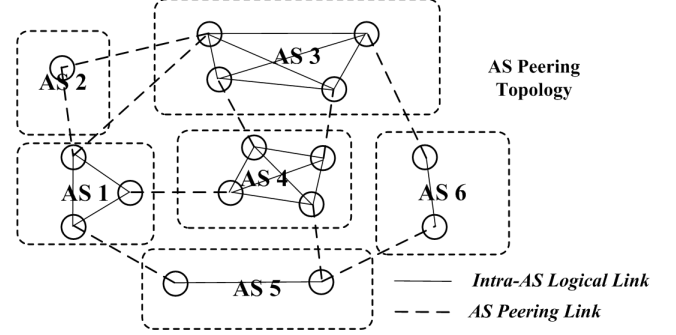


Figure 4: The corresponding AS peering level model.

We denote the AS peering system as $(\mathcal{N}, v, M_e, E_e, R_e)$, where M_e is the set of all edge routers of all ASes and E_e is the set of inter-AS peering links and logical internal links for every pair of edge routers. Figure 4 shows the corresponding AS peering topology for the network system shown in Figure 3. R_e denotes the feasible routes defined on the subgraphs of $G_e = (M_e, E_e)$. We can define R_e similarly as the feasible route R , except each intra-AS route only takes a logical link which directly connects the two edge routers of the AS. The AS peering model needs less information than the router level model, and describes different peering links between ASes. More realistically, the edge router information might be advertised by the ASes themselves, and the peering link establishments and usage can be seen via BGP routes.

The worth function v :

After introducing the AS level, the router level and the AS peering level of the network system, we construct of the worth function v for them. We define the worth function v to be the profit, i.e. the revenue minus the routing cost, as

$$v(\mathcal{S}) = v^0(\mathcal{S}) - v^c(\mathcal{S}), \quad (2)$$

where v^0 denotes the revenue and v^c denotes the routing cost. We define the revenue generated from the end-to-end service connecting AS i to AS j by W_{ij} . The revenue function v^0 is defined as the following.

$$v^0(\mathcal{S}) = \sum_{i, j \in \mathcal{S}} W_{ij} \mathbf{1}_{\{i \text{ is connected to } j \text{ by AS level graph } G_S\}}. \quad (3)$$

The revenue function depends on the topology of the network. As long as the AS i is connected to AS j , the revenue

function indicates that a revenue of W_{ij} can be obtained from the customers for providing data delivery service at the traffic intensity Λ_{ij} . We assume that a feasible route can be applied to achieve this service; however, the revenue does not depend on which route is used to achieve the service. On the other hand, the routing cost v^c not only depends on the topology of the network, but also depends on the route for achieving end-to-end services.

We start from the router level topology, defining the routing cost on each link $(i, j) \in E$. We denote c_{ij} as the routing cost function on link (i, j) , defined by $c_{ij}(x) = c_{ij}^s(x) + c_{ij}^r(x)$, where c_{ij}^s and c_{ij}^r denote the sending cost of router i and the receiving cost of router j . We assume that $c_{ij}^s(x)$ and $c_{ij}^r(x)$ are monotonically increasing with the aggregate traffic intensity x on the link and $c_{ij}^s(0) = c_{ij}^r(0) = 0$. We denote c_k as the routing cost of an AS k , defined as the aggregate sending and receiving costs of the links possessed by the AS. Given a route R on a router level topology, for any coalition \mathcal{S} , AS k 's routing cost is defined by $c_k(\mathcal{S}, R) = c_k^s(\mathcal{S}, R) + c_k^r(\mathcal{S}, R)$ where the sending cost $c_k^s(\mathcal{S}, R)$ and the receiving cost $c_k^r(\mathcal{S}, R)$ are defined as the following:

$$c_k^s(\mathcal{S}, R) = \sum_{(l_1, l_2) \in G_{\mathcal{S}}, l_1 \in m_k} c_{l_1 l_2} \left(\sum_{i, j \in M_{\mathcal{S}}} r_{ij}(\mathcal{S}, (l_1, l_2)) \right),$$

$$c_k^r(\mathcal{S}, R) = \sum_{(l_1, l_2) \in G_{\mathcal{S}}, l_2 \in m_k} c_{l_1 l_2} \left(\sum_{i, j \in M_{\mathcal{S}}} r_{ij}(\mathcal{S}, (l_1, l_2)) \right).$$

In reality, the instantaneous traffic intensity varies and the routing cost might depend on the congestion level. Here, we consider the total routing cost incurred over a certain period of time, therefore, we consider the costs as a function of the average traffic intensity. Thus, given the router level information R , we define the cost function v^c as the following:

$$v^c(\mathcal{S}, R) = \sum_{k \in \mathcal{N}} c_k(\mathcal{S}, R). \quad (4)$$

At the AS level topology, we consider each AS as a set of links and the routing costs can be imposed on each AS. Without the router level information, we assume that the cost function c_k of AS k is monotonically increasing with the traffic intensity going X_k through it. Given an AS level route $R_{\mathcal{N}}$, the routing cost function v^c becomes:

$$v^c(\mathcal{S}, R_{\mathcal{N}}) = \sum_{k \in \mathcal{N}} c_k(X_k(\mathcal{S}, R_{\mathcal{N}})), \quad (5)$$

where X_k is the traffic intensity going through AS k defined by Equation (1).

Finally, for an AS peering topology, the cost function v^c is similar to that of the router level topology. The costs can be defined on the set of links E_e , including peering links and logical intra-AS links, instead of router level links E .

3. PROFIT DISTRIBUTION MECHANISM

In this section, we formally define the class of profit distribution mechanisms ϕ and derive a specific mechanism φ based on various desirable properties. We will show that the mechanism φ is also compatible to optimal routes and smart interconnection in later sections.

Definition 2. A profit distribution mechanism is an operator ϕ on a network system (\mathcal{N}, v) that assigns a profit vector $\phi(\mathcal{N}, v) = (\phi_1, \dots, \phi_N)$ in \mathbb{R}^N . Each $\phi_i(\mathcal{N}, v)$ denotes the assigned profit of AS i .

Remark: For the network system (\mathcal{N}, v) , we do not specify if it is a router level or an AS level system. We suppose that the topology and the feasible route are fixed, therefore, all information is embedded into the worth function v defined by Equation (2). Later, when ISPs change interconnection and routing decisions, links E and routes R will appear as parameters for the network system.

3.1 Desirable Properties

We design a suitable mechanism $\phi(\mathcal{N}, v)$ that satisfies the following desirable properties among ASes.

Property 1 (EFFICIENCY). $\sum_{i \in \mathcal{N}} \phi_i(\mathcal{N}, v) = v(\mathcal{N})$. The efficiency property requires that the assigned profit balances the profit received from the service. In other words, the mechanism does not contribute or receive extra revenue. Since v is defined as the profit (revenue minus cost), all costs will be recovered from the revenue and the profit distribution mechanism determines the surplus for each AS.

Property 2 (SYMMETRY). If $v(\mathcal{S} \cup \{i\}) = v(\mathcal{S} \cup \{j\})$ for all $\mathcal{S} \in \mathcal{N} \setminus \{i, j\}$, then $\phi_i(\mathcal{N}, v) = \phi_j(\mathcal{N}, v)$.

The symmetry property requires that if two ASes contribute the same to every subset of other ASes, they should receive the same amount of profit.

Property 3 (BALANCED CONTRIBUTION). For any $i, j \in \mathcal{N}$, j 's contribution to i equals i 's contribution to j , i.e. $\phi_i(\mathcal{N}, v) - \phi_i(\mathcal{N} \setminus \{j\}, v) = \phi_j(\mathcal{N}, v) - \phi_j(\mathcal{N} \setminus \{i\}, v)$.

Here, (\mathcal{S}, v) for some $\mathcal{S} \subset \mathcal{N}$ defines the distributed profit for a sub-system of (\mathcal{N}, v) , where all ASes $\mathcal{N} \setminus \mathcal{S}$ are removed from the system and $v(\cdot)$ is restricted to the subsets of \mathcal{S} . The balanced contribution property addresses the fairness between any pair of ASes. If we start with a two-AS system $(\mathcal{N}, v) = (\{1, 2\}, v)$, the gain (or loss) from cooperation is $v(\mathcal{N}) - v(\{1\}) - v(\{2\})$. Thus, the egalitarian solution is

$$\phi_i(\mathcal{N}, v) = v(\{i\}) + \frac{1}{2}[v(\mathcal{N}) - v(\{1\}) - v(\{2\})], \quad i = 1, 2.$$

The balanced contribution property preserves and generalizes the egalitarian property in the sense that by reducing ASes recursively, the family of $\{\phi_i(\mathcal{S}, v)\}_{\mathcal{S} \subset \mathcal{N}, i \in \mathcal{S}}$ constitutes the egalitarian solutions [16].

Property 4 (DUMMY). If i is a dummy AS, i.e. $\Delta_i(v, \mathcal{S}) = 0$ for every $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$, then $\phi_i(\mathcal{N}, v) = 0$.

The dummy property requires that ASes that have no marginal contribution to any set of other ASes should receive zero profit. Because these ASes cannot improve the cooperation for making any potential profit, it's harmless to remove them from the system.

Property 5 (STRONG MONOTONICITY). *If (\mathcal{N}, v) and (\mathcal{N}, w) are two systems such that for some $i \in \mathcal{N}$, $\Delta_i(v, \mathcal{S}) \geq \Delta_i(w, \mathcal{S})$ for all $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$, then $\phi_i(\mathcal{N}, v) \geq \phi_i(\mathcal{N}, w)$.*

Property 6 (ADDITIVITY). *Given any two systems (\mathcal{N}, v) and (\mathcal{N}, w) , if $(\mathcal{N}, v + w)$ is the system where the worth function is defined by $(v + w)(\mathcal{S}) = v(\mathcal{S}) + w(\mathcal{S})$, then $\phi_i(\mathcal{N}, v + w) = \phi_i(\mathcal{N}, v) + \phi_i(\mathcal{N}, w)$ for all $i \in \mathcal{N}$.*

Both strong monotonicity and additivity properties connect the distributed profits of two systems that only differ in the worth functions. Suppose v and w represent two different types of services provided by the same group of ASes. Comparing the contribution across two different services, the strong monotonicity property requires that the more an AS contributes to a service, the more profit it receives. The additivity property requires the profit distribution mechanism to be an additive operator on the space of the worth function. Additivity property guarantees that if the service worth is additive, then the distributed profit is the sum of the profits generated by serving each individual service. In other words, profit distribution for service v will not be affected by the service w . In practice, we can consider a subset of ASes of the whole Internet which provides certain QoS service by devoting separate bandwidth provisions. The profit assigned for this QoS service only depends on the profit of the QoS service itself.

3.2 The Shapley Value Mechanism

Proposed by Lloyd Shapley [25, 20], the Shapley value is known to be the unique value satisfying all six properties above.

Definition 3. *The Shapley value φ is defined by*

$$\varphi_i(\mathcal{N}, v) = \frac{1}{N!} \sum_{\pi \in \Pi} \Delta_i(v, S(\pi, i)) \quad \forall i \in \mathcal{N}, \quad (6)$$

where Π is the set of all $N!$ orderings of \mathcal{N} and $S(\pi, i)$ is the set of players preceding i in the ordering π .

Remark: The Shapley value of an AS can be interpreted as the expected marginal contribution $\Delta_i(v, \mathcal{S})$ where \mathcal{S} is the set of ASes preceding i in a uniformly distributed random ordering. The Shapley value depends only on the values $\{v(\mathcal{S}) : \mathcal{S} \subseteq \mathcal{N}\}$.

In particular, if the routing costs are negligible compared with the revenue generated from the services, we can regard the revenue function v^0 as the worth function. In that case, the router level system $(\mathcal{N}, v, M, E, R)$, the AS level system $(\mathcal{N}, v, E_{\mathcal{N}}, R_{\mathcal{N}})$ and the AS peering level system $(\mathcal{N}, v, E_e, R_e)$ can be reduced into the system $(\mathcal{N}, v^0, E_{\mathcal{N}})$, which only depends on the structure of the AS level graph

$(\mathcal{N}, E_{\mathcal{N}})$. The Shapley values of the system $(\mathcal{N}, v^0, E_{\mathcal{N}})$ can be calculated by substituting v with v^0 in Equation (6). The value $\varphi_i(\mathcal{N}, v^0, E_{\mathcal{N}})$ is referred to as the Myerson value [17] in the literature.

4. INCENTIVE FOR OPTIMAL ROUTING

Given a fixed topology and a feasible route, the Shapley value $\varphi(\mathcal{N}, v)$ achieves various desirable properties as mentioned in last section. In this section, we still assume that the ASes form a fixed topology; however, each AS i might want to use a specific feasible route that maximizes its profit $\varphi_i(\mathcal{N}, v)$. In particular, we focus on a router level system $(\mathcal{N}, v, M, E, R)$. We drop the fixed parameters M and E , denoting the system as (\mathcal{N}, v, R) . The problem becomes that each AS i might choose a route R_i that maximizes its profit $\varphi_i(\mathcal{N}, v, R_i)$ under the Shapley value mechanism. We analyze the routes that the Shapley value mechanism induces ASes to use. Although these analysis are performed on a router level system, the results are applicable to AS level and AS peering level systems as well.

4.1 Optimal Routes and Equilibrium

From a system perspective, we wish that ASes would choose the route that maximizes the aggregate profit of the network. We define the set of all optimal routes \mathcal{R}^* as the following.

$$\mathcal{R}^* = \{R | R \in \mathcal{R}, v(\mathcal{S}, R) = \sup_{R'} v(\mathcal{S}, R') \forall \mathcal{S} \subseteq \mathcal{N}\}. \quad (7)$$

Notice that the optimal routing strategy might not be unique. We refer to R^* as any optimal route in \mathcal{R}^* .

Since the worth function $v = v^0 - v^c$, where v^0 does not depend on the route R , any optimal route $R^* \in \mathcal{R}^*$ also minimizes the aggregate routing cost $v^c(\mathcal{S})$ for any coalition \mathcal{S} . However, selfish ASes might want to minimize their own costs instead of the aggregate routing cost. Consequently, they might not follow the optimal route. We define a routing strategy R_i , a possible route chosen by AS i , as the following.

Definition 4. *A routing strategy $R_i = \{r_{jk} : j, k \in M\}$ of AS i is a set of feasible routes where each r_{jk} is defined on the domain $\{\mathcal{S} : i \in \mathcal{S} \subseteq \mathcal{N}\} \times E$.*

Notice that a routing strategy has the same definition as a feasible route R defined in Section 2.2 except it is defined on a sub-domain. R_i only contains the routes when AS i is part of the coalition \mathcal{S} . We interpret R_i as a routing strategy of AS i because it gives AS i all the possibility to change the routes when i is participating the cooperation. In reality, AS i might not be able to control all the routes for any coalition $\mathcal{S} \ni i$; however, we define a larger space of routing strategies that AS i can possibly implement.

Similarly, we denote the space of all routing strategies of AS i as \mathcal{R}_i . The set of optimal routing strategies of AS i is defined by

$$\mathcal{R}_i^* = \{R_i | R_i \in \mathcal{R}_i, v(\mathcal{S}, R) = \sup_{R'} v(\mathcal{S}, R') \forall i \in \mathcal{S} \subseteq \mathcal{N}\}. \quad (8)$$

We also refer to R_i^* as any optimal routing strategy in \mathcal{R}_i^* . Given a route R and a routing strategy R_i , we define an updated route $R \oplus R_i$ as the following.

$$R \oplus R_i(\mathcal{S}) = \begin{cases} R_i(\mathcal{S}) & \text{if } i \in \mathcal{S} \\ R(\mathcal{S}) & \text{if } i \notin \mathcal{S}. \end{cases}$$

$R \oplus R_i$ can be interpreted as the new route after AS i applies the strategy R_i to the old route R . If $R_i \in \mathcal{R}_i^*$, the $R \oplus R_i$ becomes closer to an optimal route. If each AS i applied an optimal routing strategy R_i^* sequentially on any route R , the resulting route becomes an optimal route.

The following theorem shows that under the Shapley value mechanism, each AS i can apply an optimal routing strategy R_i^* on any existing route R to maximize its own profit at $\varphi_i(\mathcal{N}, v, R \oplus R_i^*)$. Due to the space limitation, the proofs of theorems and corollaries are omitted but are available in the technical report [15].

Theorem 1 (OPTIMAL ROUTING). *Given any feasible route $R \in \mathcal{R}$, by applying an optimal routing strategy R_i^* , AS i maximizes its profit under the Shapley value mechanism, i.e. $\varphi_i(\mathcal{N}, v, R \oplus R_i^*) \geq \varphi_i(\mathcal{N}, v, R \oplus R_i)$ for all $R_i \in \mathcal{R}_i$ and $i \in \mathcal{N}$.*

Theorem 1 states that every AS can maximize its own profit by adopting an optimal routing strategy. Moreover, not only does any optimal route R^* maximize the aggregate profit $v(\mathcal{N}, R^*)$, it is also a Nash equilibrium routing strategy for all ASes.

Corollary 1. *Under the Shapley value mechanism, every optimal routing strategy $R^* \in \mathcal{R}^*$ is a Nash equilibrium.*

Corollary 1 states that when an optimal route is being used, it is compatible to all ASes' optimal routing strategies and no AS has an incentive to deviate from it. Notice that although the optimal routing strategy might not be unique, the Shapley value solution (the profit for each AS) is unique.

Theorem 2 (AS PROFIT DECOMPOSITION). *For each AS, the Shapley value profit can be decomposed into a Myerson value on the AS level topology and a Shapley value on the routing costs:*

$$\varphi_i(\mathcal{N}, v, R) = \varphi_i(\mathcal{N}, v^0, E_{\mathcal{N}}) - \varphi_i(\mathcal{N}, v^c, R) \quad \forall i \in \mathcal{N},$$

where $\varphi_i(\mathcal{N}, v^c, R)$ is the Shapley value of AS i in the system $(\mathcal{N}, v^c, M, E, R)$ that has the worth function v^c .

Theorem 2 gives a convenient way to calculate the profit for each AS by separating the AS value into an AS level Myerson value and a Shapley value restricted to the costs. More importantly, it also explains why sometimes peering links can be used to improve the aggregate profit of the system. We illustrate some examples of the AS value solutions in the next subsection.

4.2 Examples and Simulation

Figure 5(a) illustrates the first example with two coast-to-coast backbone ISPs. Customers of router 1 at the west

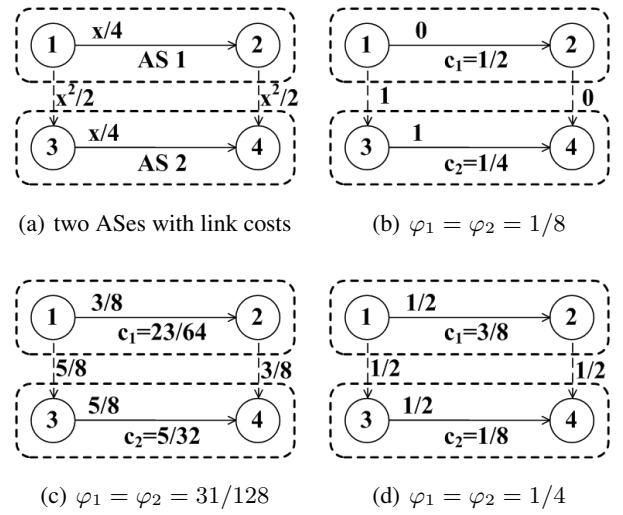


Figure 5: An example of two peering ISPs.

coast need to communicate with customers of router 4 at the east coast. We normalize the revenue and required traffic intensity to be 1. Router 1 peers with router 3 at the west coast and router 2 peers with router 4 at the east coast. We assume that all the receiving costs are zero, and the cost on a link is the same as the sending cost. The costs on intra-AS paths and inter-AS paths are $c_{12}(x) = c_{34}(x) = x/4$ (where x is the traffic carried on the link) and $c_{13}(x) = c_{24}(x) = x^2/2$. By Theorem 2, each AS obtains the same profit:

$$\begin{aligned} \varphi_i(\mathcal{N}, v, R) &= \varphi_i(\mathcal{N}, v^0, R) - \varphi_i(\mathcal{N}, v^c, R) \\ &= \frac{1}{2} - \frac{1}{2}v^c(\mathcal{N}, R) \\ &= \frac{1}{2} - \frac{1}{2}[c_1(\mathcal{N}, R) + c_2(\mathcal{N}, R)]. \end{aligned}$$

We compare the profit distributions for different routing strategies by AS 1. In Figure 5(b), AS 1 uses the hot-potato routing strategy, which routes all traffic through router 3. The routing costs of the two ASes are $c_1 = \frac{1}{2}$ and $c_2 = \frac{1}{4}$. Although AS 1 avoids using its internal link (1, 2), it does not optimize its own cost. Each AS obtains $\varphi_k = \frac{1}{8}$. In Figure 5(c), AS 1 chooses the route that minimizes its own routing cost c_1 . Both ASes improve their profit to be $\varphi_k = \frac{31}{128}$. In Figure 5(d), AS 1 uses an optimal route which minimizes aggregate routing costs for both ASes. As a result, this optimal route achieves the maximum profit $\varphi_k = \frac{1}{4}$ for both ASes, which is twice as much as the profit from hot-potato routing. Notice that no matter how much real cost an AS may carry, it will be recovered from the revenue $v^0(\mathcal{N})$. The Shapley value mechanism determines the profit of each AS from the total profit $v(\mathcal{N})$.

Figure 6 illustrates the second example where a source AS 1 wants to communicate with AS 4. Again, we normalize the revenue and required traffic intensity to be 1. Traffic must go through a core AS 3; however, AS 2 is a local peer with AS 1




Figure 6: An example of using peering link.

which can also carry traffic. We assume the sending costs on link $(3, 3a)$ and $(2, 3a)$ to be $x/2$ and x^2 respectively. We assume all other costs are negligible. The right sub-figure shows the profit distribution and the optimal routing strategy. Each AS obtains an equal profit of $\frac{9}{64}$. One way to understand this even-share solution is that any of the ASes is indispensable. For example, without AS 2, the total routing cost is 1; therefore the profit becomes zero. Theorem 2 also gives an explanation. From the AS level topology, AS 2 is a dummy AS. The Myerson values are $\varphi_2(v^0) = 0$ and $\varphi_1(v^0) = \varphi_3(v^0) = \varphi_4(v^0) = \frac{1}{3}$. However, from the cost compensation, AS 2 obtains $\varphi_2(v^c) = -\frac{9}{64}$. In this sense, we know that AS 2's profit comes from its contribution of reducing the routing cost for the end-to-end service. In general, this explains why sometimes in reality, peering links or even provider-to-customer links can also be reasonably used to provide efficiency.




Figure 7: Hot-potato routing Vs. optimal routing.

In the third example, we consider the topology in Figure 3 with six ASes. We assume each end-to-end service generates a revenue of 10, and has the required traffic intensity $\lambda_{ij} = 1$ for all pair of router i and j that do not belong to the same AS. We compare the profit distribution of the ASes under the Shapley mechanism when ASes use hot-potato routing and optimal route in Figure 7. The result confirms that the optimal routing induces more profit for all ASes than any

other non-optimal route.

4.3 Optimal Routing in Practice

The model $(\mathcal{N}, v, M, E, R)$ assumes that we know the router level topology (M, E) and the corresponding traffic intensity $\{\lambda_{ij} : i, j \in M\}$. In reality, each AS might not want to reveal its internal structure and the router-to-router traffic intensity measurement might not be feasible. However, we can still apply the optimal routing results on the corresponding AS level system $(\mathcal{N}, v, E_{\mathcal{N}}, R_{\mathcal{N}})$. The AS level topology $(\mathcal{N}, E_{\mathcal{N}})$ can be derived from ISP interconnection links and BGP routes.

Suppose each ISP i collects a total revenue of W_i from all its customers. By measuring all the traffic intensity $\{\Lambda_{ij} : j \in \mathcal{N} \setminus \{i\}\}$ from AS i to any other AS j , we can estimate the revenues W_{ij} to be $W_i \Lambda_{ij} / \sum_{k \in \mathcal{N}} \Lambda_{ik}$, assuming that the revenue is proportional to the traffic intensity directed to a destination AS j . In service contracts, the future month's required traffic intensity can be predicted and adjusted based on the historical traffic patterns between ISPs.

The optimal routing results of Theorem 1 and Corollary 1 are applicable on the AS level system $(\mathcal{N}, v, E_{\mathcal{N}}, R_{\mathcal{N}})$. In the following example, we explore Columbia University's autonomous system (AS 14) as a source ISP.




Figure 8: A snapshot of BGP routes for Columbia University on May 15, 2007.




Figure 9: Routes from Columbia to Global Crossing during May 2007.

Figure 8 shows a snapshot of the BGP routes generated by BGPPlay [1]. From time to time, the BGP paths change. We

choose the destination ISP to be the Global Crossing (AS 3549). We trace the BGP routes changes during May 2007. Figure 9 shows the active routes and the corresponding ISPs connecting Columbia University with Global Crossing. The Shapley value profits of each ISP are shown in Figure 10, assuming all ISPs have same cost functions.

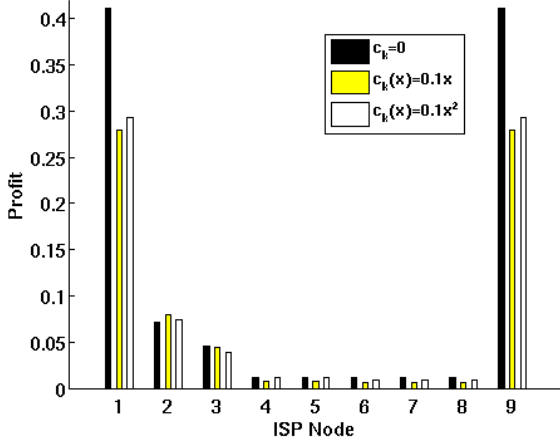


Figure 10: Revenue distribution for the ISPs.

With zero costs, each ISP obtains a Myerson value. When $c_i(x) = 0.1x$, the cost is linearly proportional to the carrying traffic and the optimal route is to use AS path $1 \rightarrow 2 \rightarrow 9$. Due to the routing cost, the sum of all profits $v(\mathcal{N})$ becomes 0.7 and most ISPs' profits decrease. However, ISP 2 (Qwest)'s profit increases from 0.072 to 0.079. This is because ISP 2 provides the optimal routing path and has a strong impact on achieving the maximum aggregate profit 0.7. When $c_i(x) = 0.1x^2$, the optimal route uses AS path $1 \rightarrow 2 \rightarrow 9$ for half of the total traffic and the three remaining AS paths for $\frac{1}{6}$ of the total traffic. The aggregate profit is improved to be $v(\mathcal{N}) = 0.75$. Since ISP 2 has a weaker impact on this solution, its profit decreases.

Because AS level topology totally ignores the internal topology of each AS, consequently, the AS level network model does not distinguish the following routing costs.

- Internal routing costs from different entering routers to different egress routers.
- Inter-AS routing costs of using parallel inter-AS links.

For example, in the topology in Figure 5, the AS level model cannot distinguish the internal costs of going through router 1 and path $1 \rightarrow 2$, as well as the two AS peering paths $1 \rightarrow 3$ and $2 \rightarrow 4$. Therefore, the AS level information is not enough to avoid selfish internal routing (e.g. hot-potato routing) for ASes.

In practice, although ASes do not reveal their internal topology very often, they export their edge routers in BGP routes. Thus, we can regard each AS i as a set of fully connected edge routers $m_i^e \subset m_i$. Only with the more delicate information of the traffic intensity of the edge routers $\{\lambda_{ij} : i, j \in$

$M_e\}$ can we apply the AS peering level model $(\mathcal{N}, v, M_e, E_e, R_e)$. In this model, we need each AS to report the true internal routing costs for each pair of its edge routers. In order to make the ASes telling the true internal routing costs, we might need some verification process when we recover each AS's real internal routing costs. With the above conditions, each AS can decide the proportion of traffic going through each inter-AS link, and tries to optimize its internal routing costs without reveal its internal topology. Notice that, if the Shapley value mechanism can be applied at this level, it will reshape the BGP inter-domain routing protocol for the ASes to cooperatively achieve an optimal route; however, keep the current intra-domain protocols compatible.

In general, the optimal routing incentive induced by the Shapley value can be shown at different levels of a network system. The more information the real network system can obtain, the more delicate level of optimal routes the ASes will be encouraged to use. Finally, we conjecture that the optimal routing practice will encourage shorter BGP paths in terms of routing costs, and diversify the usage of multiple parallel paths in routing.

5. INCENTIVE FOR INTERCONNECTING

In previous sections, we assumed that whenever a source-destination pair is connected by the graph G_S , a feasible route is performed by the coalition \mathcal{S} to achieve the end-to-end service. However, selfish ASes, whose objectives are to maximize profits, might not be willing to provide the service. For example, two ASes can provide a transit service by interconnecting with each other and obtain a revenue of w . However, it incurs a routing cost c that is larger than w . The Shapley value for each AS becomes $(w - c)/2 < 0$. It demonstrates how both ASes share the loss instead of profit. In reality, both ASes might not want to be interconnected and provide this service.

In this section, we assume that ASes are free to decide whether or not to provide an end-to-end service, as well as whether or not to interconnect with other ASes. We explore the change in the profit distribution when ASes vary the interconnection topology. We show that under the Shapley value mechanism, ASes have incentives to be well-connected so as to maximize their own profits. Like the presentation in Section 4.1, we focus on the router level network model; however, the results are general enough to be applied to both AS level and AS peering level network models.

To model the willingness of routing, we extend the domain of routes that can be used by ASes. We define an *extended route* $\tilde{R} = \{\tilde{r}_{ij}\}$, where each $\tilde{r}_{ij}(\mathcal{S})$ is either a feasible route performed by the coalition \mathcal{S} defined in Section 2.2 or a zero vector which implies that the route is not performed. We denote $\tilde{\mathcal{R}}$ as the space of all extended routes, which also includes all the feasible routes, i.e. $\mathcal{R} \subseteq \tilde{\mathcal{R}}$. With the definition of an extended route, any coalition can choose to serve certain end-to-end services and set a zero vector for other services that it does not want to provide routes.

Because when the end-to-end service is not provided, the ASes do not receive the amount of revenue from that service. We need to extend the revenue function v^0 as the following.

$$v^0(\mathcal{S}, \tilde{R}) = \sum_{i,j \in \mathcal{S}} W_{ij} \mathbf{1}_{\{i \text{ is connected to } j \text{ by } G_{\mathcal{S}} \text{ and } \tilde{R} \in \tilde{\mathcal{R}}\}}. \quad (9)$$

The worth function v defined on the extended routes becomes

$$v(\mathcal{S}, \tilde{R}) = v^0(\mathcal{S}, \tilde{R}) - v^c(\mathcal{S}, \tilde{R}), \quad (10)$$

where v^c is the same cost function defined in Equation (4). Similar to Equation (7) and (8), we define the set of optimal extended routes $\tilde{\mathcal{R}}^*$ and optimal extended routing strategies $\tilde{\mathcal{R}}_i^*$ for AS i as the following.

$$\tilde{\mathcal{R}}^* = \{\tilde{R} | \tilde{R} \in \tilde{\mathcal{R}}, v(\mathcal{S}, \tilde{R}) = \sup_{\tilde{R}'} v(\mathcal{S}, \tilde{R}') \forall \mathcal{S} \subseteq \mathcal{N}\}.$$

$$\tilde{\mathcal{R}}_i^* = \{\tilde{R}_i | \tilde{R}_i \in \tilde{\mathcal{R}}_i, v(\mathcal{S}, \tilde{R}_i) = \sup_{\tilde{R}_i'} v(\mathcal{S}, \tilde{R}_i') \forall i \in \mathcal{S} \subseteq \mathcal{N}\}.$$

We refer to \tilde{R}^* as any optimal extended route in $\tilde{\mathcal{R}}^*$ and \tilde{R}_i^* as any optimal extended routing strategy for AS i in $\tilde{\mathcal{R}}_i^*$. Notice that the extended optimal route might choose not to route for certain end-to-end services in order to maximize the worth function v .

Parallel to the optimal routing results in Section 4.1, we have the following results for extended routes \tilde{R} .

Theorem 3 (EXTENDED OPTIMAL ROUTING). *Given any extended route $\tilde{R} \in \tilde{\mathcal{R}}$, by applying an optimal routing strategy \tilde{R}_i^* , AS i maximizes its profit under the Shapley value mechanism, i.e. $\varphi_i(\mathcal{N}, v, \tilde{R} \oplus \tilde{R}_i^*) \geq \varphi_i(\mathcal{N}, v, \tilde{R} \oplus \tilde{R}_i)$ for all $\tilde{R}_i \in \tilde{\mathcal{R}}_i$ and $i \in \mathcal{N}$.*

Corollary 2. *Under the Shapley value mechanism, every optimal extended routing strategy $\tilde{R}^* \in \tilde{\mathcal{R}}^*$ is a Nash equilibrium.*

In addition, by allowing the ASes to choose whether or not to provide an end-to-end service, we guarantee that the profits of the ASes are non-negative under any route $\tilde{R}^* \in \tilde{\mathcal{R}}^*$.

Theorem 4 (NONNEGATIVITY). *$\varphi_i(\mathcal{N}, v, \tilde{R}^*) \geq 0$ for any AS $i \in \mathcal{N}$ and any optimal extended route \tilde{R}^* .*

Theorem 4 guarantees that each AS can at least recover its cost by joining the cooperation and routing traffic optimally. Notice that this result might not hold when the route is not optimal. Clearly, when an AS receives a positive profit, it has an incentive to be connected and provide the service. However, the only possible discouragement is a zero profit. The next theorem characterizes the ASes that gain zero profit.

Theorem 5. *Any AS i that has profit $\varphi_i(\mathcal{N}, v, \tilde{R}^*) = 0$ is a dummy AS, and there exists an optimal extended route $\tilde{R}'^* \in \tilde{\mathcal{R}}^*$ which does not route through AS i for all $\mathcal{S} \subseteq \mathcal{N}$.*

Theorem 5 states that if any AS receives zero profit under the Shapley value mechanism, it is a dummy AS and there is always an optimal extended route without using this AS. Consequently, although ASes that receive zero profit do not have

incentive to remain interconnected, their disconnections do not hurt the cooperation for providing services.

Interestingly, on the other hand, if an AS i does not carry any traffic in an optimal extended route with coalition \mathcal{N} , i.e. $X_i(\mathcal{N}, \tilde{R}^*) = 0$, it does not necessarily imply that AS i 's profit is zero. Because $X_i(\mathcal{S}, \tilde{R}^*)$ might be positive for some $\mathcal{S} \subset \mathcal{N}$, which means AS i provides some backup usage in case ASes $\mathcal{N} \setminus \mathcal{S}$ leave. In this case, AS i has an incentive to be interconnected and receives positive profit, although it might not actually carrying any traffic in the optimal route. The example shown in Figure 9 with cost function $c_i(x) = 0.1x$ for all ASes exhibits this situation. Although the optimal route only uses path $1 \rightarrow 2 \rightarrow 9$, AS 3 to 9 also receive positive profits.

It might be puzzling that an AS may obtain a positive profit in a system without actually carrying any traffic. However, these ASes are not dummy. They provide robustness of the network in case some of the relay ASes fail. Moreover, although these ASes share part of the total profit, they still benefit the *veto ASes* that are essential for the end-to-end services.

Definition 5. *An AS i is called a veto AS, if $v(\mathcal{S}) > 0$ for all \mathcal{S} in $\{\mathcal{S} : i \in \mathcal{S} \subseteq \mathcal{N}\}$.*

Every veto AS is essential to the end-to-end services. In other words, if any veto AS leaves the system, the service cannot be provided. In particular, for single source-destination flows, the source and destination are by nature veto ASes.

We next provide three theorems that prove the *Interconnection Incentives* of our mechanism.

Theorem 6 (MONOTONICITY – ADDING ASES). *For any veto AS i of the system $(\mathcal{N}, v, \tilde{R}^*)$, we have $\varphi_i(\mathcal{N}, v, \tilde{R}^*) \geq \varphi_i(\mathcal{S}, v, \tilde{R}^*)$ for any $\mathcal{S} \subseteq \mathcal{N}$.*

Theorem 6 tells that the full cooperation maximizes veto ASes' profits. Although some non-veto AS might not carry traffic and still obtain a positive profit, its existence still helps the cooperation and increases veto ASes' profits. Actually, a stronger statement, $\varphi_i(\mathcal{S}, v, \tilde{R}^*) \geq \varphi_i(\mathcal{T}, v, \tilde{R}^*)$ for any $\mathcal{T} \subseteq \mathcal{S} \subseteq \mathcal{N}$, can be made and the proof is similar. Theorem 6 focuses on the coalition that participates in the cooperation. The following theorems assume that the set of participating ASes is fixed. However, we explore the profits of the ASes when they decide whether or not to interconnect with neighboring ASes.

Theorem 7 (INCENTIVE FOR INTERCONNECTION). *In the system $(\mathcal{N}, v, E, \tilde{R}^*)$, suppose $l_1 \in m_i$ and $l_2 \in m_j$ are two routers belong to ASes i and j . If l_1 and l_2 are not directly connected (e.g. $(l_1, l_2) \notin E$), then adding the interconnection between l_1 and l_2 achieves no less profits for both AS i and j . Mathematically, we have $\varphi_k(\mathcal{N}, v, E \cup \{(l_1, l_2)\}, \tilde{R}^*) \geq \varphi_k(\mathcal{N}, v, E, \tilde{R}^*)$ for $k = i, j$ and any $l_1 \in m_i, l_2 \in m_j$.*

Theorem 7 addresses that by interconnecting with other ASes, one AS might be able to increase its profit. Because when an

AS connects to more ASes, it provides better robustness and connectivity for the end-to-end service. However, this might reduce other ASes' profits. The following theorem characterize the ASes whose profits are only possibly increased when more and more ASes start to interconnect.

Theorem 8 (MONOTONICITY – ADDING LINKS). *For any veto ASes i of the system $(\mathcal{N}, v, E, \tilde{R}^*)$, we have $\varphi_i(\mathcal{N}, v, E, \tilde{R}^*) \geq \varphi_i(\mathcal{N}, v, E', \tilde{R}^*)$ for any $E' \subseteq E$.* Theorem 8 states the interconnection effect to the veto ASes. When more intra-AS or inter-AS links are available for an end-to-end service, veto ASes' profit will be increased. ASes are encouraged to be interconnected by receiving a positive profit and veto ASes obtain more net revenue when the cooperation strength increases.

Remark: Theorems 7 and 8 assume that it is free to establish new links. In reality, setting up an interconnection link might induce cost to ASes. Therefore, if the extra aggregate profit (the save in the routing costs) obtained from the interconnection exceeds the cost of building the link, ASes have incentive to interconnect. This is because the costs of building the new link will be recovered from the Shapley value mechanism, and the connecting ASes would obtain more profits. Notice that, although the profits of connecting ASes and veto ASes will be increased, the total revenue paid by end users remain unchanged. Under the Shapley value mechanism, ASes have incentives to interconnect so as to reduce routing costs and maximize their own profits.

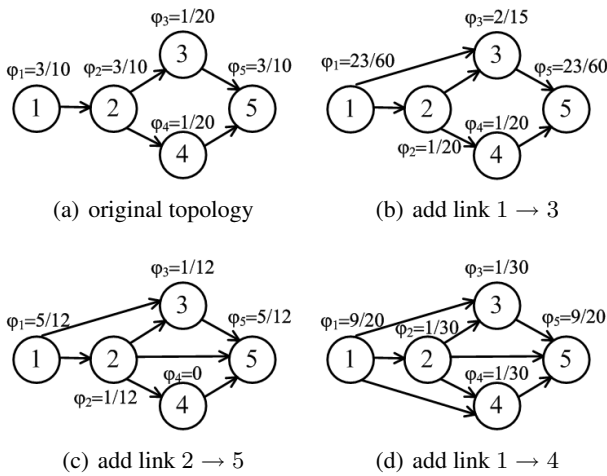


Figure 11: Monotonicity of veto ASes when adding links.

Figure 11 illustrates the changes in profit distribution when ASes start to interconnect with neighboring ASes. We ignore the routing costs and focus on an AS level topology. In Figure 11(a), AS 1, 2 and 5 are veto ASes. In Figure 11(b), link $1 \rightarrow 3$ is added. AS 2 is no longer a veto AS and its value decreases; however, AS 1 and 5's value increase. AS 3's value also increases since its direct connection with the source provides robustness. In Figure 11(c), link $2 \rightarrow 5$ is further added. As a result, AS 4 becomes a dummy AS, and

again the veto AS 1 and 5's values are increased. Similarly, AS 2's value increases as its direct connection with the destination provides robustness. In Figure 11(d), link $1 \rightarrow 4$ is added. After directly connecting to the source AS, AS 4 becomes a parallel AS to 2 and 3 and is no longer dummy. Notice that after this topology is created, links $2 \rightarrow 3$ and $2 \rightarrow 4$ become dummy and might not be used.

6. RELATED WORK

The study of Internet interconnection started a decade ago. Srinagesh [23] studied the cost structures of various ISPs and their consequences in interconnection agreements. Both Bailey [3] and Huston [12] surveyed the existing interconnection settlements. Huston [12] and Frieden [8, 9] also compared the existing Internet settlement models with that of the telecommunication industry's. Bailey concluded that bilateral agreements might be suitable for large ISPs while cooperative agreements might work for small ones. Huston concluded that the zero-dollar peering and the customer/provider relationships were the only stable models for the Internet at the time. Gao [11] proposed a relationship-based model for ISPs and categorized the interconnection relationship by provider-to-customer, peer-to-peer and sibling-to-sibling links. Instead of modeling bilateral relationships of ISPs, our work models the cooperations among multiple ISPs as a whole and designs a multilateral settlement for all ISPs to share profits.

Roughgarden et al. [21] analyzed the performance degeneration caused by selfish routing in terms of latency. Teixeira et al. [24] conducted experiments and found that hot-potato routing causes longer delays and slow convergence for BGP routes. Johari et al. [14] showed that hot-potato routing could be three times more expensive than optimal routing. Feigenbaum et al. [7] used mechanism design [18] approaches to encourage ASes to use minimum cost routes. This approach operates in the way that the source and the destination ASes want to optimize a "supply chain" for routing. Our approach, however, treats each AS equally and divides the total profit fairly among a team of collaborators.

Frieden [8, 9] discussed the consequence of Internet Balkanization: Interconnection has begun to shift from a widespread, voluntary and non-discriminatory model to a hierarchical and discriminatory model; and ISPs currently avoid the burdens of common carriage. Network neutrality [5, 26, 10] proponents criticized the discriminatory behavior by ISPs, believing that it harms the productivity, innovation and end-to-end connectivity of the Internet. However, most of the network neutrality debate has been focused on the potential regulatory enforcements, by which telephony companies have been regulated. Wu [26] surveyed the discriminatory practices of broadband provider and cable operators, and proposed solutions of bandwidth management and policing for ISPs to avoid *broadband discrimination*. Nonetheless, few work has been done on network neutrality on peering agreement. Crowcroft [5] reviewed technical aspects of network neutrality and concluded that we should not engineer

for network neutrality. Like Wu's proposal for broadband providers, our work proposes a profit distribution mechanism for ISPs. Without re-engineering for the network neutrality, this approach encourages ISPs to interconnect and alleviates the discriminatory interconnection problem.

Game theory [19, 16] has been applied to different network areas. Mostly, non-cooperative games [22, 4] have been used to model the selfish behaviors of network entities. Our work incorporates the Shapley value solution from *coalition games* [19, 6, 13] to model the cooperative nature of the ISPs. Different from non-cooperative games, coalition game does not specify the minute description of individual players, e.g. the strategies, order of move and corresponding payoff consequences. Instead, coalition game reduces all information into the possible profits generated by each coalition. As mentioned by Eyal Winter in [25], the major advantage of this approach is its practical usefulness in a multi-player environment, which provides a more tractable structure than non-cooperative games.

7. CONCLUSION

In this paper, we propose a novel multilateral settlement for ISPs. Under this multilateral settlement, customers pay for the end-to-end services provided by a set of ISPs, and ISPs collectively share the revenue generated from these customers based on a profit distribution mechanism. We design a profit distribution mechanism that can be applied for network systems with different levels of information: AS level, AS peering level and router level systems. The profit distribution mechanism implements the Shapley value solution, which satisfies efficiency and various fairness properties. More importantly, we show that under the Shapley value mechanism, selfish ISPs have incentives to adopt global optimal routing strategies instead of local greedy ones, as well as to interconnect with neighboring ISPs so as to maximize their own profits. In particular, we prove that not only do the global optimal routes maximize the aggregate profit of the network system, they are also Nash equilibrium solutions for all ISPs to follow. In addition, locally connecting to more neighboring ISPs will increase an ISP's profit. As a result, veto ISPs' profits will be monotonically increasing under the Shapley value mechanism when interconnections become more prevalent. Finally, in order to enforce the proposed profit distribution mechanism, future directions of this work should include the consideration of timescale, granularity, and trust issues of the network protocol which implements this mechanism.

8. REFERENCES

- [1] BGPplay from Route Views.
- [2] Federal Communications Commission (FCC) – Telecommunications Act of 1996. [Online] Available: <http://www.fcc.gov/Reports/tcom1996.pdf>.
- [3] J. P. Bailey. The economics of Internet interconnection agreements. *Internet economics*, pages 155 – 168, 1997.
- [4] X. Cao, H. Shen, R. Milito, and P. Wirth. Internet pricing with a game theoretical approach: concepts and examples. *IEEE/ACM Transactions on Networking*, 10(2):208–216, April 2002.
- [5] J. Crowcroft. Net neutrality: the technical side of the debate: a white paper. *ACM SIGCOMM Computer Communication Review*, 37(1), January 2007.
- [6] G. Demange and M. Wooders. *Group formation in economics: networks, clubs, and coalitions*. Cambridge University Press, Cambridge, 2005.
- [7] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. In *Proceedings of the 2002 ACM Symposium on Principles of Distributed Computing*, 2002.
- [8] R. Frieden. Without public peer: the potential regulatory and universal service consequences of Internet Balkanization. *Virginia Journal of Law and Technology*, 3, 1998.
- [9] R. Frieden. When internet peers become customers: the consequences of settlement-based interconnection, 2000.
- [10] R. Frieden. Network neutrality or bias? – handicapping the odds for a tiered and branded Internet, 2006.
- [11] L. Gao. On inferring autonomous system relationships in the internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, December 2001.
- [12] G. Huston. *ISP Survival Guide: Strategies for Running a Competitive ISP*. John Wiley and Son, New York, 1999.
- [13] M. O. Jackson. Allocation rules for network games. *Game Theory and Information from EconWPA*, 2003.
- [14] R. Johari and J. Tsitsiklis. Routing and peering in a competitive internet. *43rd IEEE Conference on Decision and Control*, 2:1556–1561, 2004.
- [15] R. T. B. Ma, D. M. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. The Shapley value mechanism for ISP settlement. *technical report, Columbia University, EE Department, July 2007*. [Online] Available: <http://dnawsl.cs.columbia.edu/pubsdb/citation/paperfile/145/tech07.pdf>.
- [16] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- [17] R. Myerson. Graphs and cooperation in games. *Mathematics of Operations Research*, 2:225–229, 1977.
- [18] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proc. 31st Annual Symposium on Theory of Computing (STOC99)*, 1999.
- [19] M. J. Osborne and A. Rubinstein. *A course in game theory*. The MIT Press Cambridge, Massachusetts, 1994.
- [20] A. Roth. *The Shapley value: Essays in honor of Lloyd S. Shapley*. Cambridge University Press, Cambridge, 1988.
- [21] T. Roughgarden and E. Tardos. How bad is selfish routing? In *IEEE Symposium on Foundations of Computer Science*, pages 93–102, 2000.
- [22] S. Shakkottai and R. Srikant. Economics of network pricing with multiple ISPs. In *Proceedings of IEEE Infocom*, 2005.
- [23] P. Srinagesh. Internet cost structures and interconnection agreements. *Internet economics*, pages 121 – 154, 1997.
- [24] R. Teixeira, A. Shaikh, T. Griffin, and J. Rexford. Dynamics of hot-potato routing in IP networks. *Proceedings of the ACM SIGMETRICS/Performance*, 2004.
- [25] E. Winter. *The Shapley Value, in The Handbook of Game Theory*. R. J. Aumann and S. Hart, North-Holland, 2002.
- [26] T. Wu. Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law*, 141, 2005.