Internet Economics: The Use of Shapley Value for ISP Settlement

Richard T. B. Ma, Dah Ming Chiu, *Fellow, IEEE*, John C. S. Lui, *Fellow, IEEE*, Vishal Misra, *Member, IEEE*, and Dan Rubenstein, *Member, IEEE*

***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 and encourage ISP connectivity so as to limit balkanization.**

***Index Terms*—Coalition game, incentives, ISP settlement, Nash equilibrium, Shapley value.**

I. Introduction

THE Internet is composed of thousands of connected *autonomous systems* (ASs). Before transitioning to the private sector, these ASs' primary focus was to improve connectivity and network performance—who got paid was not the primary concern. However, in its current form, ISPs, each composed of one or more ASs, have a primary interest to maximize their own profit. Connectivity is currently implemented via bilateral contracts that are generally either a peering relationship where ISPs offer to carry one another's traffic or a customer–provider relationship where one ISP pays the other for transit [15].

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 [26]. 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 [11]. 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], [30]. 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 properties that formally indicate that an ISP's profit is proportional to its contribution to the value of the network. Specifically, our contributions are the following.

- We propose (Section II) a novel multilateral settlement model, where customers pay for end-to-end services and ISPs collectively share the revenue for providing services.
- We implement this settlement via a mechanism based on the Shapley value and show that the following are achieved (Section III).
  - **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, fairness, 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

Manuscript received July 30, 2008; revised May 06, 2009; approved by IEEE/ACM Transactions on Networking Editor J. Walrand. First published May 17, 2010; current version published June 16, 2010. A conference version of this paper [18] appeared in CoNEXT'07.

R. T. B. Ma is with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: tbma@ee.columbia.edu).

D. M. Chiu is with the Department of Information Engineering, The Chinese University of Hong Kong, Shatin NT, Hong Kong (e-mail: dmchiu@ie.cuhk.edu.hk).

J. C. S. Lui is with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Shatin NT, Hong Kong (e-mail: cslui@cse.cuhk.edu.hk).

V. Misra and D. Rubenstein are with the Department of Computer Science, Columbia University, New York, NY 10027 USA (e-mail: misra@cs.columbia.edu; danr@cs.columbia.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2010.2049205

TABLE I
MAJOR NOTATION USED IN THIS PAPER

$\mathcal{N}$	set of all Autonomous Systems
$\mathcal{S}$	any subset of ASes
M	set of all routers
E	set of directed inter-router links
$\hat{E}$	set of directed inter-AS links
λ_{ij}	required traffic rate from router i to router j
Λ_{ij}	required traffic rate from AS i to AS j
W_{ij}	revenue of end-to-end service from AS i to AS j
R	the flow profile used by ASes to route traffic
R_i	routing strategy of AS i to route traffic
R_i^*	optimal routing strategy of AS i to route traffic
$\tilde{R}$	the extended flow profile used by ASes to route traffic
$\tilde{R}_i$	extended routing strategy of AS i to route traffic
$\tilde{R}_i^*$	optimal extended routing strategy of AS i to route traffic
v	worth function (profit) defined on any set of ASes
v^0	revenue component of the worth function
v^c	cost component of the worth function
C_k	routing cost of AS k
φ_i	the Shapley value (profit) of AS i

topology converges to a Nash equilibrium where the aggregate profit of the system is also maximized (Section IV).

— **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 V).

- We illustrate examples of profit distribution among ISPs, including the real AS topology of Columbia University, New York, NY. We also simulate and compare the profits under hot-potato routing and optimal routing generated by our new mechanism. We show that by using the new mechanism, every ISP increases profits.

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: 1) Our profit redistribution is *not a zero sum game*, and in fact all participating players stand to gain; and 2) 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. To aid our discussion, Table I summarizes the major notation that we use in this paper.

II. COOPERATIVE FRAMEWORK

A. 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 Fig. 1. At the bottom

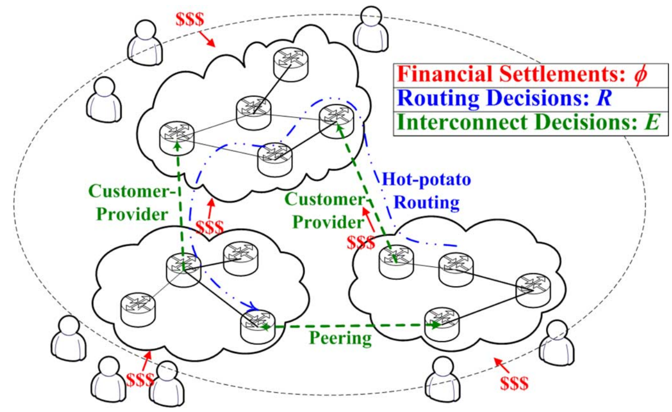

Fig. 1. A view of ISP interactions of the Internet.

layer, pairs of ISPs decide whether or not to 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 its customer ISPs forward traffic.

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 routing is often used to choose the closest egress point based on the intradomain 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 one another, 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.

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 Fig. 2, as the following.

S1) Customers make service agreements (charge and requirements) at their local ISPs; however, from the 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 routing R to maximize its profit.

In the first stage, the service negotiation at the edge of the network does not require users to buy resources from multiple ISPs along the communication paths. 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 their origins or destinations. Commercial content providers might be charged more than nonprofit 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 mech-*

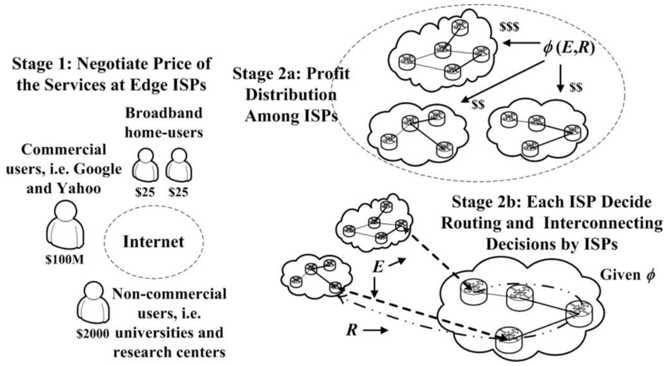

Fig. 2. A two-stage multilateral settlement model.

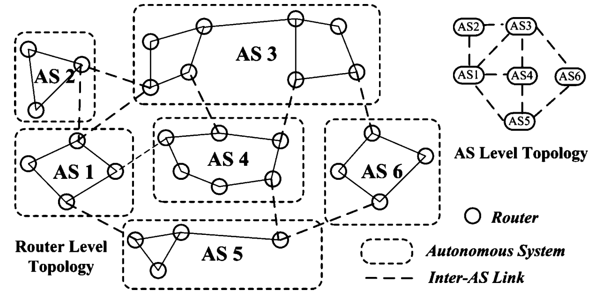

Fig. 3. A router-level model and its corresponding AS-level topology.

anism $\phi(E, R)$ that encourages selfish ISPs to interconnect extensively and route efficiently.

B. Network Model

We first define the network flows that will be used in our network models as follows.

Definition 1: A *flow* f on a directed graph $G = (V, E)$ is a mapping function $f : E \rightarrow \mathbb{R}$.

Remark: Each $f((i, j))$ can be considered as the traffic rate going from vertex i to j through the link $(i, j) \in E$ on graph G. Suppose we want to achieve end-to-end data delivery rates $\{\lambda_{ij} : i, j \in V\}$, we define the flows that achieve these rates.

Definition 2: A *single feasible flow* f_{ij} is a flow that achieves a source–destination rate λ_{ij} from vertex i to vertex j on a directed graph $G = (V, E)$.

Remark: If i and j are not connected in graph G, by default we define f_{ij} by a zero vector. Mathematically, each nonzero vector f_{ij} satisfies the following flow conservation constraints:

$$\sum_{(i,k_1)\in E} f_{ij}\left((i,k_1)\right) = \sum_{(k_2,j)\in E} f_{ij}\left((k_2,j)\right) = \lambda_{ij}.$$
$$\sum_{(k_1,l)\in E} f_{ij}\left((k_1,l)\right) = \sum_{(l,k_2)\in E} f_{ij}\left((l,k_2)\right) \quad \forall l \in V \setminus \{i,j\}.$$

Definition 3: A *feasible flow* f is a flow that achieves rates $\{\lambda_{ij} : i, j \in V\}$ on a directed graph G. It can be written as a summation of single feasible flows $f = \sum_{i,j\in V} f_{ij}$.

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

To fairly distribute profits among ASs, 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 ASs. We call any nonempty subset $\mathcal{S} \subseteq \mathcal{N}$ a *coalition* of the ASs. Each coalition can be thought of as a subnetwork 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 subnetworks formed by all coalitions. In other words, for any coalition $\mathcal{S}$, $v(\mathcal{S})$ defines the profit generated by the subnetwork formed by the set of ASs $\mathcal{S}$. Through the worth function v, we can measure the contribution of an AS to a group of ASs as the following.

Definition 4: 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 ASs. In the rest of this section, we develop our network model and the corresponding worth function.

Router-Level Model: We denote $(\mathcal{N}, v, M, E, R)$ as the router-level network system. M denotes 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 $i \neq j$. We denote $G_{\mathcal{S}}$ as the subgraph of G induced by $\mathcal{S}$, defined by $G_{\mathcal{S}} = (M_{\mathcal{S}}, E_{\mathcal{S}})$, where $M_{\mathcal{S}} = \bigcup_{i\in\mathcal{S}} m_i$ and $E_{\mathcal{S}} = \{(i,j) \in E : i, j \in M_{\mathcal{S}}\}$. $G_{\mathcal{S}}$ is the router-level topology formed by the coalition $\mathcal{S}$. We assume that the network needs to provide end-to-end data delivery services to customers, and these services require the network to achieve router-to-router traffic rates $\{\lambda_{ij} \geq 0 : i, j \in M\}$. R denotes a *flow profile* used by the ASs to route traffic at these rates.

Definition 5: A *flow profile* R maps each coalition $\mathcal{S} \subseteq \mathcal{N}$ to a feasible flow $f_{\mathcal{S}}$ on the induced directed subgraph $G_{\mathcal{S}}$.

Remark: Each $R(\mathcal{S})$ is a feasible flow defined on the subgraph $G_{\mathcal{S}}$, which can be considered as a routing strategy used by coalition $\mathcal{S}$ to route data traffic.

Given the router-level topology $G = (M, E)$, we can construct the corresponding AS-level topology $\hat{G} = (\mathcal{N}, \hat{E})$. $\hat{E}$ denotes the set of directed logical links between the ASs, defined by $\hat{E} = \{(i,j) | (k,l) \in E, k \in m_i, l \in m_j\}$. Fig. 3 illustrates an example of a router-level topology and the corresponding AS-level topology. At the AS level, the subgraph $\hat{G}_{\mathcal{S}}$, formed by the coalition $\mathcal{S}$, is defined by $\hat{G}_{\mathcal{S}} = (\mathcal{S}, \hat{E}_{\mathcal{S}})$, where $\hat{E}_{\mathcal{S}} = \{(i,j) \in \hat{E} : i, j \in \mathcal{S}\}$.

Worth Function v: We define the worth function v to be the profit, i.e., the revenue minus the routing cost, as

$$v(\mathcal{S}) = v^0(\mathcal{S}) - v^c(\mathcal{S}) \qquad (1)$$

where v^0 denotes the revenue and v^c denotes the routing cost. We say that AS i is *connected* to AS j by $\hat{G}_{\mathcal{S}}$ if there is some $k \geq 1$ and a sequence $(n_0, n_1, \ldots, n_k)$ such that $n_0 = i$, $n_k =$

j, and $(n_{l-1}, n_l) \in \hat{E}_{\mathcal{S}}$ for $1 \le l \le k$. 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.

$$v^0(\mathcal{S}) = \sum_{i,j \in \mathcal{S}} W_{ij} \mathbf{1}_{\{i \text{ is connected to } j \text{ by AS level graph } \hat{G}_{\mathcal{S}}\}}. \tag{2}$$

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 services. We assume that a feasible flow is used to achieve this service; however, the revenue does not depend on which flow 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 flows that achieve data delivery end-to-end services.

We denote c_{ij} as the routing cost function on link $(i,j) \in E$, 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 any feasible flow $f_{\mathcal{S}}$ defined on the subgraph $G_{\mathcal{S}}$ of a router-level topology, AS k's routing cost $C_k(f_{\mathcal{S}})$ is defined by

$$C_k(f_{\mathcal{S}}) = \sum_{\substack{(i,j) \in G_{\mathcal{S}} \\ i \in m_k}} c_{ij}^s \left(f_{\mathcal{S}}\left((i,j)\right)\right) + \sum_{\substack{(i,j) \in G_{\mathcal{S}} \\ j \in m_k}} c_{ij}^r \left(f_{\mathcal{S}}\left((i,j)\right)\right).$$

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 any flow profile R, we define the cost function v^c as the following:

$$v^c(\mathcal{S}) = \sum_{k \in \mathcal{S}} C_k\left(R(\mathcal{S})\right). \tag{3}$$

III. 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. In Sections IV and V, we will show that the mechanism φ is also compatible to optimal routing and smart interconnection.

Definition 6: A *profit distribution mechanism* is an operator ϕ on a network system $(\mathcal{N}, v)$ that assigns a profit vector $\phi(\mathcal{N}, v) = (\phi_1, \ldots, \phi_N)$ 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 suppose that the topology and the flow profile are fixed, therefore all information is embedded into the worth function v defined by (1). Later, when ISPs change interconnection and routing decisions, links E and flow profile R will appear as parameters for the network system.

A. Desirable Properties

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

Property 1 (Efficiency): $\sum_{i \in \mathcal{N}} \phi_i(\mathcal{N}, v) = v(\mathcal{N})$.

The efficiency property requires that the profit assigned equals 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. If we focus on the system that consists of both the revenue-generating entities, i.e., ISPs, and the revenue contributing entities, i.e., customers, the efficiency property is equivalent to the *budget balance* [21] condition in the literature of *mechanism design* [23], under which no money runs out of or into the system.

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 ASs contribute the same to every subset of other ASs, they should receive the same amount of profit.

Property 3 (Fairness): 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 subsystem of $(\mathcal{N}, v)$, where all ASs $\mathcal{N} \setminus \mathcal{S}$ are removed from the system and $v(\cdot)$ is restricted to the subsets of $\mathcal{S}$. The fairness property addresses the fairness between any pair of ASs. 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

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

The fairness property preserves and generalizes the egalitarian property in the sense that by reducing ASs recursively, the family of $\{\phi_i(\mathcal{S}, v)\}_{\mathcal{S} \subset \mathcal{N}, i \in \mathcal{S}}$ constitutes the egalitarian solutions [21].

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 ASs that have no marginal contribution to any other coalitions should receive zero profit. Because these ASs cannot contribute for making any potential profit, it is 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 $i \in \mathcal{N}$, $\Delta_i(v, \mathcal{S}) \ge \Delta_i(w, \mathcal{S})$ for all $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$, then $\phi_i(\mathcal{N}, v) \ge \phi_i(\mathcal{N}, w)$.

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 ASs. 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 that the profit distribution mechanism is an additive operator on the space of the worth function. The additivity property guarantees that if the worth function of the service is additive, then the distributed profit is the sum of the profits generated by serving each individual service. In other words, the profit distribution for service v will not be affected by the service w. In practice, we can consider a subset of ASs of the Internet that provides certain QoS service by devoting separate bandwidth provisions. The assigned profit associated with the QoS service will only depend on the profit generated by the QoS service itself.

B. Shapley Value Mechanism

Proposed by Lloyd Shapley [25], [30], the Shapley value is the unique value that satisfies all six properties above.

Definition 7: The *Shapley value* φ is defined by

$$\varphi_i(\mathcal{N}, v) = \frac{1}{N!}\sum_{\pi\in\Pi}\Delta_i\left(v, S(\pi, i)\right) \quad \forall i \in \mathcal{N} \tag{4}$$

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 ASs preceding i in a uniformly distributed random ordering.

In particular, if the routing costs are negligible when compared to the revenue generated from the services, we can regard the revenue function v^0 as the worth function. In that case, the network system $(\mathcal{N}, v, M, E, R)$ can be reduced to the system $(\mathcal{N}, v^0, \hat{E})$, which only depends on the structure of the AS-level topology $(\mathcal{N}, \hat{E})$. The Shapley values of the system $(\mathcal{N}, v^0, \hat{E})$ can be calculated by substituting v^0 for v in (4). The value $\varphi_i(\mathcal{N}, v^0, \hat{E})$ is referred to as the Myerson value [22] in the literature, where the worth function only depends on the interconnecting topology.

IV. Incentive for Optimal Routing

Given a fixed topology and a flow profile, the Shapley value $\varphi(\mathcal{N}, v)$ achieves various desirable properties as mentioned in the last section. In this section, we still assume that the ASs form a fixed topology; however, each AS i might want to use a specific flow profile that maximizes its profit $\varphi_i(\mathcal{N}, v)$. We consider a network system $(\mathcal{N}, v, M, E, R)$, which we drop the fixed parameters M and E, and refer to it as $(\mathcal{N}, v, R)$. The problem becomes that each AS i might choose a flow profile R_i that maximizes its profit $\varphi_i(\mathcal{N}, v, R_i)$ under the Shapley value mechanism. We analyze the flows that the Shapley value mechanism induces ASs to use.

A. Optimal Flow and Equilibrium

From a global perspective, we wish that ASs would choose the feasible flows that maximize the aggregate profit of the network. We define $\mathcal{R}$ as the space of all flow profiles R. We define the set of optimal flow profiles $\mathcal{R}^*$ as the following:

$$\mathcal{R}^* = \left\{R | R \in \mathcal{R}, v(\mathcal{S}, R) = \sup_{R'} v(\mathcal{S}, R') \quad \forall \mathcal{S} \subseteq N\right\}. \tag{5}$$

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

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

Definition 8: A *routing strategy* R_i of AS i maps each coalition $\mathcal{S}$ that AS i belongs to, i.e., $i \in \mathcal{S} \subseteq \mathcal{N}$, to a feasible flow $f_{\mathcal{S}}$ on the induced directed subgraph $G_{\mathcal{S}}$.

Notice that a routing strategy has the same definition as a flow profile R defined in Section II-B, except it is defined on a subdomain. R_i only contains the feasible flows carried by coalitions $\mathcal{S}$ that AS i belongs to. We interpret R_i as a routing strategy of AS i because it gives AS i the possibility of changing the flows when i is participating in the coalition. In reality, AS i might not be able to control all the flows 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

$$\mathcal{R}_i^* = \left\{R_i | R_i \in \mathcal{R}_i, v(\mathcal{S}, R) = \sup_{R'} v(\mathcal{S}, R') \quad \forall i \in \mathcal{S} \subseteq N\right\}. \tag{6}$$

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

$$R \oplus R_i(\mathcal{S}) = \begin{cases} R_i(\mathcal{S}), & \text{if } i \in \mathcal{S} \\ R(\mathcal{S}), & \text{if } i \notin \mathcal{S}. \end{cases} \tag{7}$$

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

The following theorem shows that under the Shapley value mechanism, each AS i can apply an optimal routing strategy R_i^* on any existing flow profile R to maximize its own profit at $\varphi_i(\mathcal{N}, v, R \oplus R_i^*)$.

Theorem 1 (Optimal Routing): Given any flow profile $R \in \mathcal{R}$, 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 $i \in \mathcal{N}$.

Proof: We compare the marginal contribution of AS i to any coalition $\mathcal{S}$ that it does not belong to, i.e., $i \notin \mathcal{S}$, under two systems that use the flow profiles R and $R \oplus R_i^*$ respectively.

$$\begin{aligned} &\Delta_i\left(v\left(\mathcal{S}, R \oplus R_i^*\right), \mathcal{S}\right) \\ &\quad = v\left(\mathcal{S} \cup \{i\}, R \oplus R_i^*\right) - v\left(\mathcal{S}, R \oplus R_i^*\right) \\ &\quad = v\left(\mathcal{S} \cup \{i\}, R \oplus R_i^*\right) - v(\mathcal{S}, R) \\ &\quad \geq v\left(\mathcal{S} \cup \{i\}, R\right) - v(\mathcal{S}, R) \\ &\quad = \Delta_i\left(v(\mathcal{S}, R), \mathcal{S}\right). \end{aligned}$$

The second equality holds because $R \oplus R_i^*(\mathcal{S}) = R(\mathcal{S})$ for any $\mathcal{S}$ that $i \notin \mathcal{S}$ by definition in (7). The inequality holds because $R \oplus R_i^*(\mathcal{S} \cup \{i\}) = R_i^*(\mathcal{S} \cup \{i\})$, and the optimal strategy R_i^* achieves no less *worth* for the coalition $\mathcal{S} \cup \{i\}$. Therefore, $\Delta_i(v(\mathcal{S}, R \oplus R_i^*), \mathcal{S}) \geq \Delta_i(v(\mathcal{S}, R), \mathcal{S}) \forall \mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$. By the strong monotonicity property, we conclude that $\varphi_i(\mathcal{N}, v, R \oplus R_i^*) \geq \varphi_i(\mathcal{N}, v, R \oplus R_i)$. ■

Theorem 1 states that every AS can maximize its own profit by adopting an optimal routing strategy. Intuitively, this is because when an AS uses an optimal routing strategy, the profits of the coalitions that this AS belongs to increase. Since the Shapley value is a weighted sum of all marginal contributions that this AS contributes to different coalitions, an AS can increase its profit-share by using an optimal routing strategy accordingly. Moreover, not only does any optimal flow profile R^* maximize the aggregate profit $v(\mathcal{N}, R^*)$, it is also a Nash equilibrium for all ASs.

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

Proof: Given any optimal flow profile R^*, an AS i can always deviate from it and apply a routing strategy R_i. The new flow profile becomes $R^* \oplus R_i$. However, there exists an $R_i^* \in \mathcal{R}_i^*$ such that after applying it to the updated flow profile, the profile changes back to R^*, i.e., $R^* \oplus R_i \oplus R_i^* = R^*$. This routing strategy is the reverse action that AS i takes to restore the flow profile back to R^*. Suppose an optimal routing strategy R^* is not a Nash equilibrium. Then, there exist an AS i with a routing strategy R_i, which $\varphi_i(\mathcal{N}, v, R^* \oplus R_i) > \varphi_i(\mathcal{N}, v, R^*)$ satisfies. However, since $R^* \in \mathcal{R}_i^*$ as well, the result contradicts Theorem 1. ■

Corollary 1 states that when an optimal flow profile is being used, it is compatible to all ASs' 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 (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 cost

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

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.

Proof: Since the worth function can be written as $v = v^0 - v^c$, by the additivity property, the value $\varphi_i(\mathcal{N}, v, R)$ becomes

$$\begin{aligned} \varphi_i(\mathcal{N}, v, R) &= \varphi_i(\mathcal{N}, v^0 - v^c, R) \\ &= \varphi_i(\mathcal{N}, v^0, R) - \varphi_i(\mathcal{N}, v^c, R) \\ &= \varphi_i(\mathcal{N}, v^0, \hat{E}) - \varphi_i(\mathcal{N}, v^c, R). \end{aligned}$$

The last equality holds because v^0 is insensitive to the routing costs, and the Shapley value is equivalent to the Myerson value introduced in Section III-B. ■

In defining the revenue function v^0 in (2), we assumed that when a coalition forms a connected graph, all end-to-end data delivery services are provided and the corresponding revenues are generated. Under this assumption, Theorem 2 gives a convenient way to decompose the Shapley profit into an AS-level Myerson value revenue and a router-level Shapley value cost. In general, connected ASs may refuse to provide services and generate less than a full amount of revenue. However, the Shapley profit can always be decomposed as a Shapley revenue component minus a Shapley cost component, where the revenue component does not only depend on the AS topology and therefore cannot be simply represented by the Myerson value on the topology. Practically, this result also explains why sometimes inter-AS 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.

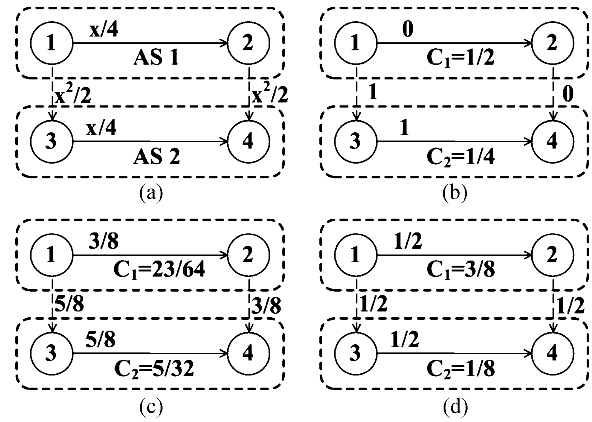

Fig. 4. An example of two backbone ISPs. (a) Two ASs with link costs. (b) $\varphi_1 = \varphi_2 = 1/8$. (c) $\varphi_1 = \varphi_2 = 31/128$. (d) $\varphi_1 = \varphi_2 = 1/4$.

B. Examples and Simulation

Fig. 4(a) illustrates a first example with two coast-to-coast backbone ISPs. Customers of router 1 on the west coast need to communicate with customers of router 4 on the east coast. We normalize the revenue and required traffic intensity to be 1. Router 1 peers with router 3 on the west coast, and router 2 peers with router 4 on 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

$$\begin{aligned} \varphi_i(\mathcal{N}, v, R) &= \varphi_i(\mathcal{N}, v^0, \hat{E}) - \varphi_i(\mathcal{N}, v^c, R) \\ &= \frac{1}{2} - \frac{1}{2} v^c(\mathcal{N}) \\ &= \frac{1}{2} - \frac{1}{2} \left[C_1 \left(R(\mathcal{N}) \right) + C_2 \left(R(\mathcal{N}) \right) \right]. \end{aligned}$$

We compare the profit distributions for different routing strategies by AS 1. In Fig. 4(b), AS 1 uses the hot-potato routing strategy, which routes all traffic through router 3. The routing costs of the two ASs are $C_1 = 1/2$ and $C_2 = 1/4$. Although AS 1 avoids using its internal link $1 \rightarrow 2$, it does not optimize its own cost. Each AS obtains $\varphi_k = 1/8$. In Fig. 4(c), AS 1 chooses the feasible flow that minimizes its own routing cost c_1. Both ASs improve their profit to be $\varphi_k = 31/128$. In Fig. 4(d), AS 1 uses an optimal flow that minimizes aggregate routing costs for both ASs. As a result, this optimal flow achieves the maximum profit $\varphi_k = 1/4$ for both ASs, 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 $v(\mathcal{N})$. Fig. 5 illustrates a second example where a source AS 1 wants to communicate with AS 4. Again,

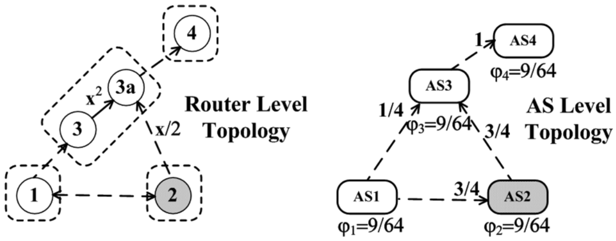

Fig. 5. An example of using a peering link.

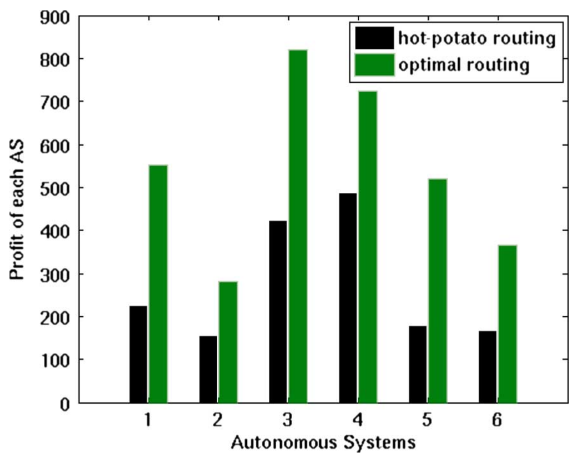

Fig. 6. Hot-potato routing versus optimal routing.

we normalize the revenue and required traffic intensity to be 1. Traffic must go through AS 3; however, AS 2 is a local peer with AS 1, which can also carry traffic. We assume the sending costs on link $3 \rightarrow 3a$ and $2 \rightarrow 3a$ to be $x/2$ and x^2, respectively. We assume all other costs are negligible. The right subfigure shows the profit distribution and the optimal routing strategy. Each AS obtains an equal profit of 9/64. One way to understand this even-share solution is that any of the ASs 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 $\varphi_1(v^0) = \varphi_3(v^0) = \varphi_4(v^0) = 1/3$. However, from the cost compensation, AS 2 obtains $\varphi_2(v^c) = -(9/64)$. In this sense, we know that AS 2's profit comes from its contribution of reducing the routing cost. In general, this explains why sometimes in reality, peering links or even provider-to-customer links are used to provide efficiency. In a third example, we consider the topology of six ASs shown in Fig. 3. We assume each end-to-end service generates a revenue of 10 and has the required traffic intensity $\lambda_{ij} = 1$ for all pairs of routers i and j that do not belong to the same AS. We compare the profit distribution of the ASs under the Shapley mechanism when ASs use hot-potato routing and optimal routing in Fig. 6. The result confirms that optimal routing induces more profit for all ASs than any other nonoptimal routing strategies.

V. Incentive for Interconnecting

In previous sections, we assumed that whenever a source–destination pair is connected by the graph $G_{\mathcal{S}}$, the coalition $\mathcal{S}$ achieves an end-to-end data delivery service using a feasible flow $R(\mathcal{S})$. However, selfish ASs, whose objectives are to maximize profits, might not be willing to provide the service. For example, two ASs 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 ASs share the loss instead of profit. In reality, both ASs might not want to be interconnected and provide the service.

In this section, we assume that ASs are free to decide whether or not to provide an end-to-end data delivery service, as well as whether or not to interconnect with other ASs. We explore the change in the profit distribution when ASs vary the interconnection topology. We show that under the Shapley value mechanism, ASs have incentives to be well connected so as to maximize their own profits.

To model the willingness of routing, we define an *extended flow profile* as follows.

Definition 9: An *extended flow profile* $\tilde{R}$ maps each coalition $\mathcal{S} \subseteq \mathcal{N}$ to a feasible flow $f_{\mathcal{S}}$ on the induced directed subgraph $G_{\mathcal{S}}$ or a zero vector.

Remark: This definition extends the domain of flow profiles that can be used by ASs. Each $\tilde{R}(\mathcal{S})$ denotes either a feasible flow carried by the coalition $\mathcal{S}$ defined in Section II-B or a zero vector, which implies that no flow is carried.

We define $\tilde{\mathcal{R}}$ to be the space of all extended flow profiles, which also includes all the flow profiles, i.e., $\mathcal{R} \subseteq \tilde{R}$. With the definition of an extended flow profile, 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. Now, the revenue generated by a coalition $\mathcal{S}$ depends on the end-to-end data delivery services that it provides. We need to extend the revenue function v^0 to the following:

$$v^0(\mathcal{S}, \tilde{R}) = \sum_{i,j \in \mathcal{S}} W_{ij} \mathbf{1}_{\{i \text{ is connected to } j \text{ by } \hat{G}_{\mathcal{S}} \text{ and } \tilde{R} \in \mathcal{R}\}}. \quad (8)$$

The worth function v defined on the extended flows becomes

$$v(\mathcal{S}, \tilde{R}) = v^0(\mathcal{S}, \tilde{R}) - v^c(\mathcal{S}, \tilde{R}) \quad (9)$$

where v^c is the same cost function defined in (3).

To extend the set of optimal flow profiles in (5) and the set of optimal routing strategies in (6), we define the set of optimal extended flow profiles $\tilde{\mathcal{R}}^*$ and optimal extended routing strategies $\tilde{\mathcal{R}}_i^*$ for AS i as the following:

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

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

We refer to $\tilde{R}^*$ as any optimal extended flow profile 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 flow profile might choose not to carry flows for certain end-to-end services in order to maximize the worth function v.

Parallel to the optimal routing results in Section IV-A, we have the following results for extended flow profile $\tilde{R}$.

Theorem 3 (Extended Optimal Routing): Given any extended flow profile $\tilde{R} \in \tilde{\mathcal{R}}$, 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 $i \in \mathcal{N}$.

Proof: The same arguments from Theorem 1 apply. ■

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

Proof: The same arguments from Corollary 1 apply. ■

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

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

Proof: Since the Shapley value is a linear combination (with positive coefficients) of marginal contribution $\Delta_i(v, \mathcal{S})$, we prove that $\Delta_i(v, \mathcal{S}) = v(\mathcal{S} \cup \{i\}, \tilde{R}^*) - v(\mathcal{S}, \tilde{R}^*) \geq 0$ for all $\mathcal{S} \subset \mathcal{N}$. Since the optimal extended flow profile $\tilde{R}^*$ is being used, both $v(\mathcal{S} \cup \{i\}, \tilde{R}^*)$ and $v(\mathcal{S}, \tilde{R}^*)$ are nonnegative. If $v(\mathcal{S}, \tilde{R}^*) = 0$, the result is trivial. If $v(\mathcal{S}, \tilde{R}^*) > 0$, then the flow $\tilde{R}^*(\mathcal{S})$ is not a zero vector. We conclude that $v(\mathcal{S} \cup \{i\}, \tilde{R}^*) \geq v(\mathcal{S}, \tilde{R}^*)$ because adding AS i can only reduce the transit cost under the optimal routing decision $\tilde{R}^*$. ■

Theorem 4 guarantees that each AS can at least recover its cost by joining the whole coalition $\mathcal{N}$ and routing traffic optimally. Notice that this result might not hold when the routing strategy 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 ASs 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 flow profile $\tilde{R'}^* \in \tilde{\mathcal{R}}^*$ that does not route through AS i for all $\mathcal{S} \subseteq \mathcal{N}$.

Proof: As shown in the proof of Theorem 4, the marginal contributions are nonnegative, i.e., $\Delta_i(v, \mathcal{S}) \geq 0$ for all $\mathcal{S} \subset \mathcal{N}$. Since $\varphi_i(\mathcal{N}, v, \tilde{R}^*) = 0$, we have $\Delta_i(v, \mathcal{S}) = 0$ for all $\mathcal{S} \subset \mathcal{N}$. Hence, we know $v(\mathcal{S} \cup \{i\}, \tilde{R}^*) = v(\mathcal{S}, \tilde{R}^*)$. This implies that when AS i joins the coalition $\mathcal{S}$, it does not improve the worth. Thus, we can always apply an optimal extended flow $\tilde{R'}^*(\mathcal{S})$ to $\tilde{R'}^*(\mathcal{S} \cup \{i\})$ and assign zero throughput for AS i. ■

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 flow without using this AS. Consequently, although ASs 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 flow $\tilde{R}^*(\mathcal{N})$ with the whole coalition $\mathcal{N}$, it does not necessarily imply that AS i's profit is zero. Because AS i might carry traffic in an optimal flow $\tilde{R}^*(\mathcal{S})$ for some $\mathcal{S} \subset \mathcal{N}$, which means AS i provides some backup usage in case ASs $\mathcal{N} \setminus \mathcal{S}$ leave. In this case, AS i has an incentive to be interconnected and receives positive profit, although it might not actually be carrying any traffic in the optimal flow for the whole coalition $\mathcal{N}$. A later example shown in Fig. 10 (with cost function $C_i(x) = 0.1x$) exhibits this situation. Although the optimal flow only uses path $1 \to 2 \to 9$, AS 3 to 8 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 ASs are not dummy, and they provide robustness of the network in case some of the relay ASs fail. Moreover, although these ASs share part of the total profit, they still benefit the *veto ASs* that are essential for generating profit.

Definition 10: 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 in order to generate the profit $v(\mathcal{S})$ for any coalition. In other words, if any veto AS leaves the system, no positive profit can be generated by any coalition that does not contain the veto AS. In particular, for single source–destination flows, the source and destination are by nature veto ASs.

Next, we provide three theorems that prove the *Interconnection Incentives* of our mechanism.

Theorem 6 (Monotonicity—Adding ASs): 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 $\mathcal{S} \subseteq \mathcal{N}$.

Proof: We first consider the case where $\mathcal{S} = \mathcal{N} \setminus \{j\}$ for some AS $j \in \mathcal{N}$. When j does not connect, it simply obtains zero profit. By the fairness property of the Shapley value, we have $\varphi_i(\mathcal{N}, v, \tilde{R}^*) - \varphi_i(\mathcal{N} \setminus \{j\}, v, \tilde{R}^*) = \varphi_j(\mathcal{N}, v, \tilde{R}^*) - \varphi_j(\mathcal{N} \setminus \{i\}, v, \tilde{R}^*)$. Because i is a veto AS, $\varphi_j(\mathcal{N} \setminus \{i\}, v, \tilde{R}^*) = 0$. By Theorem 4, $\varphi_j(\mathcal{N}, v, \tilde{R}^*) \geq 0$. As a result, $\varphi_i(\mathcal{N}, v, \tilde{R}^*) \geq \varphi_i(\mathcal{N} \setminus \{j\}, v, \tilde{R}^*)$. Finally, we can successively reduce ASs from $\mathcal{N}$ to reach arbitrary coalition $\mathcal{S}$, and the monotonicity also holds. ■

Theorem 6 tells us that the whole coalition maximizes veto ASs' profits. Although some nonveto AS might not carry traffic and still obtain a positive profit, its existence still helps the cooperation and increases veto ASs' 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 change of size of a cooperative coalition. The following theorems assume a fixed cooperative coalition of ASs. However, we explore the profits of the ASs when they decide whether or not to interconnect with neighboring ASs.

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 belonging to ASs 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 lesser 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 $l_1 \in m_i, l_2 \in m_j$.

Proof: Let $\tilde{E} = E \cup \{(i, j)\}$. Since the set of links E and $\tilde{E}$ are different for the two systems, we define $v^*(\mathcal{S}, E) = v(\mathcal{S}, E, \tilde{R}^*)$ to be the worth function applied to the topology $G = (\mathcal{N}, E)$ with optimal routing strategy $\tilde{R}^*$ for coalition $\mathcal{S} \subseteq \mathcal{N}$. For $k = i$ or j, we have

$$\begin{aligned}
&\varphi_k(\mathcal{N}, v, \tilde{E}, \tilde{R}^*) \\
&\quad = \frac{1}{|\mathcal{N}|!} \sum_{\mathcal{S} \subseteq \mathcal{N} \setminus \{k\}} \left[v^*\left(\mathcal{S} \cup \{k\}, \tilde{E}\right) - v^*(\mathcal{S}, \tilde{E}) \right] \\
&\quad = \frac{1}{|\mathcal{N}|!} \sum_{\mathcal{S} \subseteq \mathcal{N} \setminus \{k\}} \left[v^*\left(\mathcal{S} \cup \{k\}, \tilde{E}\right) - v^*(\mathcal{S}, E) \right].
\end{aligned}$$

The last equation holds because when either i or j is not in the coalition $\mathcal{S}$, E and $\tilde{E}$ gives the same topology for the induced graph $G_{\mathcal{S}}$. By subtracting two values, we have

$$\varphi_k(\mathcal{N}, v, \tilde{E}, \tilde{R}^*) - \varphi_k(\mathcal{N}, v, E, \tilde{R}^*) = \frac{1}{|\mathcal{N}|!} \sum_{\mathcal{S} \subseteq \mathcal{N} \setminus \{k\}} \left[v^*\left(\mathcal{S} \cup \{k\}, \tilde{E}\right) - v^*\left(\mathcal{S} \cup \{k\}, E\right) \right].$$

Then, it is enough to show that $v^*(\mathcal{S} \cup \{k\}, \tilde{E}) \geq v^*(\mathcal{S} \cup \{k\}, E)$ for all $\mathcal{S} \subseteq \mathcal{N}$. Since $E \subseteq \tilde{E}$, the induced graph $G_{\mathcal{S}}$ is also included in $\tilde{G}_{\mathcal{S}}$ for any coalition $\mathcal{S}$. Therefore, since the optimal routing is used, $v^*(\mathcal{S}, \tilde{E})$ can achieve at least as much as $v^*(\mathcal{S}, E)$. ■

Theorem 7 states that by interconnecting with other ASs, one AS might be able to increase its profit. This is because when an AS connects to more ASs, it provides better robustness and connectivity for providing end-to-end services. However, this result does not imply that all pairs of ASs should be connected for the following two reasons. First, redundant links that do not reduce the total routing cost will not increase profits for the interconnecting ASs. In this case, ASs' profits will remain the same. Second, our model does not consider the capital expenditure of establishing a new link; however, we can extend our model to include this factor. Intuitively, if the savings on global routing cost are greater than the fixed cost of building a new link, ASs will still benefit from interconnecting with each other. In general, ASs' profits are affected by other ASs' interconnecting decisions. The following theorem characterizes the ASs whose profits might possibly be increased as more ASs connect.

Theorem 8 (Monotonicity—Adding Links): For any veto ASs 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 $E' \subseteq E$.

Proof: We define $v^*(\mathcal{S}, E) = v(\mathcal{S}, E, \tilde{R}^*)$. Since i is a veto AS, $v(\mathcal{S}) = 0$ for all $\mathcal{S}$ with $i \notin \mathcal{S}$. We have

$$\varphi_i(\mathcal{N}, v, E, \tilde{R}^*) = \frac{1}{|\mathcal{N}|!} \sum_{\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}} \Delta_i(v, \mathcal{S}) = \frac{1}{|\mathcal{N}|!} \sum_{\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}} v^*\left(\mathcal{S} \cup \{i\}, E\right).$$

Then, it is enough to show that $v^*(\mathcal{S} \cup \{i\}, E) \geq v^*(\mathcal{S} \cup \{i\}, E')$ for all $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$. Since $E' \subseteq E$, the induced graph $G'_{\mathcal{S}}$ is also included in $G_{\mathcal{S}}$ for any coalition $\mathcal{S}$. Therefore, since the optimal routing is used, $v^*(\mathcal{S}, E)$ can achieve at least as much as $v^*(\mathcal{S}, E')$. ■

Theorem 8 states the interconnection effect to the veto ASs. When more intra-AS or inter-AS links are available, veto ASs' profits will increase. ASs are encouraged to be connected by receiving more profit, and veto ASs will also benefit when the coalition becomes more robust.

Remark: Theorems 7 and 8 assume that establishing new links do not induce fixed costs. In reality, setting up an interconnection link might induce cost to ASs. Therefore, if the extra aggregate profit (due to savings in the routing cost) obtained from the interconnection exceeds the cost of building the link, ASs have an incentive to interconnect. This is because the costs of building the new link will be recovered from the Shapley value mechanism, and the connecting ASs would obtain more profits.

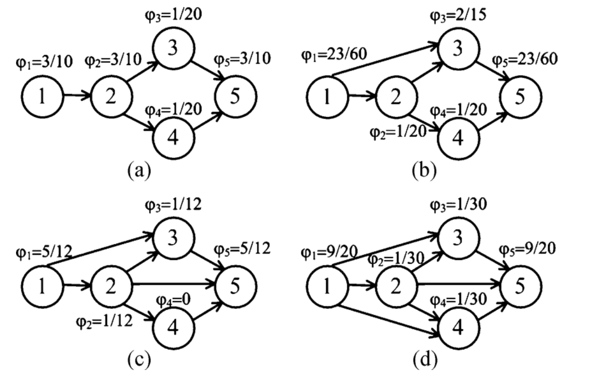

Fig. 7. Monotonicity of veto ASs when adding links. (a) Original topology. (b) Add link $1 \to 3$. (c) Add link $2 \to 5$. (d) Add link $1 \to 4$.

Notice that although the profits of connecting ASs and veto ASs will increase, the total revenue paid by end-users remain unchanged. Under the Shapley value mechanism, ASs have incentives to interconnect so as to reduce routing costs and maximize their own profits.

Fig. 7 illustrates the changes in profit distribution when ASs start to interconnect with neighboring ASs. We ignore the routing costs and focus on an AS-level topology. In Fig. 7(a), ASs 1, 2, and 5 are veto ASs. In Fig. 7(b), link $1 \to 3$ is added. AS 2 is no longer a veto AS, and its profit decreases; however, ASs 1 and 5's profits increase. AS 3's profit also increases since its direct connection with the source provides robustness. In Fig. 7(c), link $2 \to 5$ is further added. As a result, AS 4 becomes a dummy AS, and again the veto ASs 1 and 5's profits increase. Similarly, AS 2's profit increases as its direct connection with the destination provides robustness. In Fig. 7(d), link $1 \to 4$ is added. After directly connecting to the source, AS 4 becomes a parallel AS to 2 and 3 and is no longer dummy. Under this topology, links $2 \to 3$ and $2 \to 4$ become dummy and might not be used.

VI. Implementation Issues

We have shown the routing and interconnecting incentive properties of the Shapley-value profit distribution mechanism in the last two sections. In this section, we address various implementation issues so as to achieve the Shapley-value profit-share in practice.

A. Traffic and Topology Information

Our network 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, ASs might not want to reveal their internal structures, and router-to-router traffic intensity measurements might not be feasible. However, our results are applicable to network systems with different levels of information as well, e.g., AS level and AS peering level, as we describe below.

AS-Level Information: Without knowing the router-level information, we can derive the AS-level topology $\hat{G} = (\mathcal{N}, \hat{E})$ from inter-AS links and BGP routes. The required AS-to-AS traffic intensity $\{\Lambda_{ij} : i, j \in \mathcal{N}\}$ is just the aggregate router-level traffic intensity, i.e., $\Lambda_{ij} = \sum_{k \in m_i, l \in m_j} \lambda_{kl} \forall i, j \in \mathcal{N}$. Each Λ_{ij} can be measured as the volume of traffic contracted

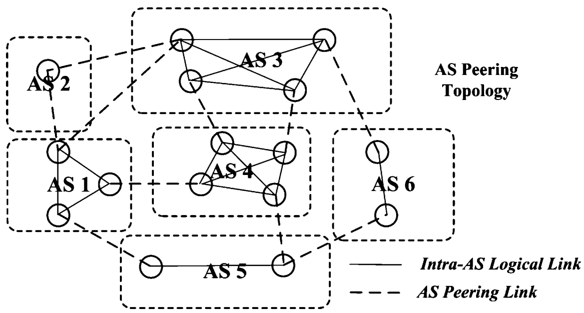

Fig. 8. The corresponding AS peering-level model.

over a certain period of time for delivery. In the current Internet [9], the timescale for these contracts is often monthly. By a similarly definition, we can also define the flow profile on the AS-level topologies $\hat{G}_S$ that achieves AS-to-AS traffic rates $\{\Lambda_{ij} \geq 0 : i, j \in \mathcal{N}\}$.

AS Peering-Level Information: Because AS-level information does not distinguish the multiple peering points between two ASs, it is often too coarse to accurately describe the network system. In the AS peering-level model, each AS does not need to reveal its internal router-level topology, but 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 routers forms a logical, fully connected graph. Fig. 8 shows the corresponding AS peering topology for the network system shown in Fig. 3. The AS peering model needs less information than what is required at the router level and describes different peering links between ASs. In practice, the edge router information is advertised by the ASs themselves, and the peering link establishments and usage can be observed via BGP routes.

In general, the routing and interconnecting incentives induced by the Shapley value can be shown at different levels of a network system. The more information the real system can obtain, the more delicate level of optimal routing and interconnecting strategies the ASs will be encouraged to use.

B. Computational Complexity

The calculation of the Shapley value is known to be $\#P$-hard [3]. In a follow-up paper [20], we discuss how to calculate the Shapley profit-share for various types of ISPs. We derive closed-form solutions for ISPs under regular topologies as well as via a dynamic programming procedure that calculates solutions under general topologies. In some situations, suitable bilateral contracts can implement the Shapley-value result. For example, in [19], we derive the bilateral payments between ISPs that can achieve the Shapley-value solution under a "content-eyeball" model. Another paper [6] also examines the bilateral prices that can achieve the Shapley-value solution in ISP peering. However, current bilateral practices often reach a solution that largely deviates from the Shapley-value solution because the Shapley solution implies that two unconventional bilateral settlements—i.e., "reverse customer–provider" and "paid peering" [20]—are needed to achieve solutions that are close to the Shapley-value solution.

C. Truth Revealing and Adaptation

In the Shapley-value mechanism, we assume that the worth of each coalition is known. However, this assumption requires knowledge of the true revenue and cost of each ISP: information that ISPs might not want to reveal. In order to maximize profit, they might even have incentives to misreport their private information. Two possible solutions are the following. First, it may be necessary to establish some regulatory policy that requires ISPs to report true revenue and cost, which are subject to audit by the government. Governmental intervention is often needed to settle ISP disputes, e.g., Level 3's depeering with Cogent in 2005. Second, it may be necessary to design a new level of "truth revealing" mechanism, which distributes profit not only based on the Shapley value, but also based on ISPs' reported information. The concept of this kind of "incentive compatible" mechanism was first introduced by Hurwicz [14]. Although the "incentive compatible" structure of the Shapley value mechanism is a future direction for our research, we believe that a tradeoff between incentive compatibility and efficiency may achieve this goal.

In either of the above solutions, we require a central mediator to collect some global information for the Shapley value mechanism as inputs. It can be a governmental institution or an association of ISPs. Given a topology, one can store the information of marginal impacts among the ASs. When existing ASs leave or new ASs join, one can calculate the Shapley-value profit by a dynamic programming procedure [20], which utilizes the previous marginal impacts and updates these information. After topological changes, each AS wants to adapt to a new optimal flow; however, the Shapley-value mechanism does not restrict ASs to use optimal flows. ASs can probe optimal flows gradually. Even if some AS cannot find the optimal flow, or be irrational to use a suboptimal flow, other ASs can still adjust their flows to minimize global routing cost. Thus, the Shapley-value mechanism is very robust in the sense that some ASs' irrational behavior will not have significant negative impact on the system.

D. Example of Optimal Routing in Practice

Suppose each ISP i collects a total revenue of W_i from all of its customers. By measuring all the traffic intensity $\{\Lambda_{ij} : j \in \mathcal{N} \setminus \{i\}\}$ from AS i to any other AS j, one 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 to the AS-level system $(\mathcal{N}, v, \hat{E})$. In the following example, we explore Columbia University's autonomous system (AS 14) as a source ISP.

Fig. 9 shows a snapshot of the BGP routes generated by BGPlay [1]. From time to time, the BGP paths change. We choose the destination ISP to be the Global Crossing (AS 3549) and trace the BGP routes changes during May 2007. Fig. 10 shows the active routes and the corresponding ISPs connecting Columbia University with Global Crossing. The Shapley-value profits of each ISP are shown in Fig. 11, assuming all ISPs have the same cost function. The cost function C_k of AS k

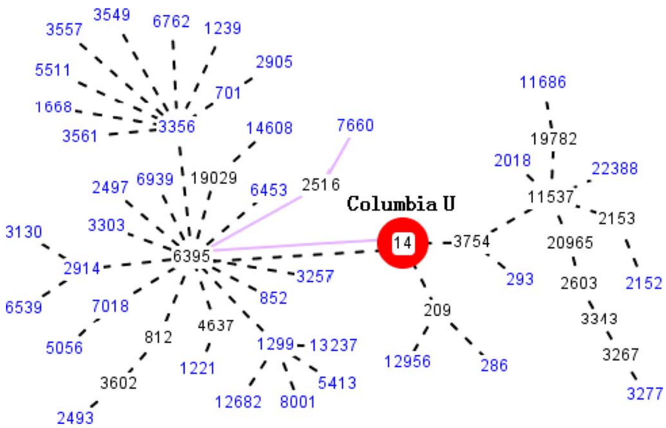

Fig. 9. A snapshot of BGP routes for Columbia University on May 15, 2007.

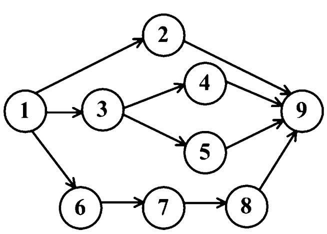

Node	ISP	AS Number
1	Columbia U	14
2	Qwest	209
3	Broadwing	6395
4	SprintLink	1239
5	Level3	3356
6	NYSERNET	3754
7	Internet2	11537
8	NTT America	3948 & 2914
9	Global Crossing	3549

Fig. 10. Routes from Columbia to Global Crossing during May 2007.

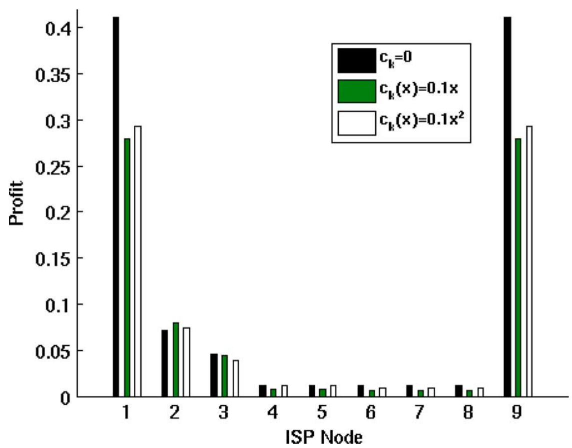

Fig. 11. Revenue distribution for the ISPs.

is monotonically increasing with the traffic intensity going through it.

With no costs, each ISP obtains a Myerson value. When $C_k(x) = 0.1x$, the cost is linearly proportional to the carrying traffic and the optimal flow uses 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 is crucial to the maximum aggregate profit of 0.7. When $C_k(x) = 0.1x^2$, the optimal flow uses AS path $1 \rightarrow 2 \rightarrow 9$ for half of the total traffic and the three remaining AS paths for 1/6 of the total traffic. The aggregate profit is improved to $v(\mathcal{N}) = 0.75$. Since ISP 2 is less crucial to 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 ingress routers to different egress routers;
- inter-AS routing costs of using parallel inter-AS links.

For example, in the topology in Fig. 4, the AS-level model cannot distinguish the internal costs of going through router 1 and path $1 \rightarrow 2$ from the two AS peering paths $1 \rightarrow 3$ and $2 \rightarrow 4$. Thus, the AS-level information is not enough to avoid ASs' selfish internal routing (e.g., hot-potato routing).

In practice, although ASs do not reveal their internal topology very often, they export their edge routers in BGP routes. Thus, one can treat each AS i as a set of fully connected edge router IDs as shown in Fig. 8. In this AS peering model, each AS must report the true internal routing costs for each pair of its edge routers. In order to make the ASs tell the true internal routing costs, we 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 try to optimize its internal routing costs without revealing its internal topology. Notice that if the Shapley value mechanism can be applied at this level, it will reshape the BGP interdomain routing protocol for the ASs to cooperatively achieve an optimal flow; however, keep the current intradomain protocols unchanged. Thus, 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.

VII. Related Work

Many research interests have been focusing on the Internet interconnections. Srinagesh [28] studied the cost structures of various ISPs and their consequences in interconnection agreements. Both Bailey [4] and Huston [15] surveyed the existing interconnection settlements. Huston [15] and Frieden [11] also compared existing Internet settlement models with those of the telecommunication industry's. Bailey concluded that bilateral agreements are suitable for large ISPs while cooperative agreements 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 [13] 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.* [26] analyzed the performance degeneration caused by selfish routing in terms of latency. Teixeira *et al.* [29] conducted experiments and found that hot-potato routing causes longer delays and slow convergence for BGP routes. Johari *et al.* [17] showed that hot-potato routing could be three times more expensive than optimal routing. Feigenbaum *et al.* [10] used mechanism design [23] approaches to encourage ASs to use minimum-cost flows. This approach operates in the way that the source and the destination ASs want to optimize a "supply chain" for routing. Our approach, however, treats each AS equally and divides profit fairly among a team of collaborators.

Frieden [11] discussed the consequence of Internet Balkanization: The process of ISP interconnection has begun to shift from a widespread, voluntary, and nondiscriminatory model to a hierarchical and discriminatory model, and ISPs currently avoid the burdens of a common carrier. Network neutrality [7], [12], [31] proponents criticize 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 focuses on the potential regulatory enforcements, by which telephony companies have been regulated. Wu [31] surveys the discriminatory practices of broadband provider and cable operators and proposes solutions of bandwidth management and policing for ISPs to avoid *broadband discrimination*. Nonetheless, little work has been done on the ISP settlement aspect of network neutrality. Crowcroft [7] reviews technical aspects of network neutrality and concludes that one should not engineer for network neutrality. Like Wu's proposal for broadband providers, our work proposes a profit distribution mechanism for ISPs. Without reengineering for network neutrality, this approach encourages ISPs to interconnect and alleviates the discriminatory interconnection problem.

Game theory [21], [24] has been applied to different network areas. Mostly, noncooperative games [5], [27] are used to model the selfish behaviors of network entities. Our work incorporates the Shapley-value solution from *coalition games* [8], [16], [24] to model the cooperative nature of the ISPs. Different from noncooperative games, a 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 Winter in [30], the major advantage of this approach is its practical usefulness in a multiplayer environment, which provides a more tractable structure than noncooperative games.

VIII. Conclusion

In this paper, we propose a novel multilateral settlement for ISPs. Under this multilateral settlement, customers pay for 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 routing strategies 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 protocols that implement this mechanism.

References

[1] BGPlay from Route Views.

[2] Federal Communications Commission (FCC), "Telecommunications Act of 1996," [Online]. Available: http://www.fcc.gov/Reports/tcom1996.pdf

[3] Y. Bachrach, E. Markakis, A. D. Procaccia, J. S. Rosenschein, and A. Saberi, "Approximating power indices," in *Proc. 7th Int. Joint Conf. Autonom. Agents Multiagent Syst.*, 2008, pp. 943–950.

[4] J. P. Bailey, "The economics of Internet interconnection agreements," *Internet Econ.*, pp. 155–168, 1997.

[5] X. Cao, H. Shen, R. Milito, and P. Wirth, "Internet pricing with a game theoretical approach: Concepts and examples," *IEEE/ACM Trans. Netw.*, vol. 10, no. 2, pp. 208–216, Apr. 2002.

[6] Y. Cheung, D. Chiu, and J. Huang, "Can bilateral ISP peering lead to network-wide cooperative settlement," in *Proc. 17th ICCCN*, 2008, pp. 1–6.

[7] J. Crowcroft, "Net neutrality: The technical side of the debate: A white paper," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 1, pp. 49–56, 2007.

[8] G. Demange and M. Wooders*, Group Formation in Economics: Networks, Clubs, and Coalitions*. Cambridge, U.K.: Cambridge Univ. Press, 2005.

[9] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr, "Complexity of Internet interconnections: Technology, incentives and implications for policy," in *Proc. 35th Res. Conf. Commun., Inf. Internet Policy (TPRC)*, 2007.

[10] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker, "A BGP-based mechanism for lowest-cost routing," in *Proc. ACM Symp. Principles Distrib. Comput.*, 2002, pp. 173–182.

[11] R. Frieden, "Without public peer: The potential regulatory and universal service consequences of Internet Balkanization," *Virginia J. Law Technol.*, vol. 3, 1998.

[12] R. Frieden, "Network neutrality or bias?—Handicapping the odds for a tiered and branded Internet," 2006.

[13] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, Dec. 2001.

[14] L. Hurwicz, "On informationally decentralized systems," in *Decision and Organization*, C. B. Maguire and R. Radner, Eds. New York: Elsevier, 1972.

[15] G. Huston*, ISP Survival Guide: Stratagies for Running a Competitive ISP*. New York: Wiley, 1999.

[16] M. O. Jackson, "Allocation rules for network games," *Games Econ. Behav.*, vol. 5, no. 1, pp. 128–154, Apr. 2005.

[17] R. Johari and J. Tsitsiklis, "Routing and peering in a competitive Internet," in *Proc. 43rd IEEE Conf. Decision Control*, 2004, vol. 2, pp. 1556–1561.

[18] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "Internet economics: The use of Shapley value for ISP settlement," in *Proc. ACM CoNEXT*, Dec. 2007, Article no. 6.

[19] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "Interconnecting eyeballs to content: A Shapley value perspective on ISP peering and settlement," in *Proc. ACM NetEcon*, Aug. 2008, pp. 61–66.

[20] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "On cooperative settlement between content, transit and eyeball Internet service providers," in *Proc. ACM CoNEXT*, Dec. 2008, Article no. 7.

[21] A. Mas-Colell, M. D. Whinston, and J. R. Green*, Microeconomic Theory*. Oxford, U.K.: Oxford Univ. Press, 1995.

[22] R. Myerson, "Graphs and cooperation in games," *Math. Oper. Res.*, vol. 2, pp. 225–229, 1977.

[23] N. Nisan and A. Ronen, "Algorithmic mechanism design," in *Proc. 31st Annu. SSTOC*, 1999, pp. 129–140.

[24] M. J. Osborne and A. Rubinstein*, A Course in Game Theory*. Cambridge, MA: MIT Press, 1994.

[25] A. Roth*, The Shapley Value: Essays in Honor of Lloyd S. Shapley*. Cambridge, U.K.: Cambridge Univ. Press, 1988.

[26] T. Roughgarden and E. Tardos, "How bad is selfish routing?," in *Proc. IEEE Symp. Found. Comput. Sci.*, 2000, pp. 93–102.

[27] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple ISPs," in *Proc. IEEE INFOCOM*, 2005, vol. 1, pp. 184–194.

[28] P. Srinagesh, "Internet cost structures and interconnection agreements," *Internet Econ.*, pp. 121–154, 1997.
[29] R. Teixeira, A. Shaikh, T. Griffin, and J. Rexford, "Dynamics of hot-potato routing in IP networks," in *Proc. ACM SIGMETRICS/Perform.*, 2004, pp. 307–319.
[30] E. Winter, *The Shapley Value, in The Handbook of Game Theory*, R. J. Aumann and S. Hart, Eds. Amsterdam, The Netherlands: North-Holland, 2002.
[31] T. Wu, "Network neutrality, broadband discrimination," *J. Telecommun. High Technol. Law*, vol. 2, pp. 141–179, 2003.

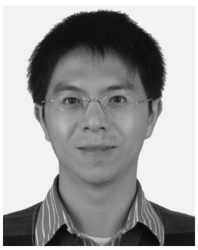

Richard T. B. Ma received the B.Sc. degree (first-class honors) in computer science and the M.Phil. degree in computer science and engineering from the Chinese University of Hong Kong in July 2002 and July 2004, respectively. He is currently working toward the Ph.D. degree at Columbia University, New York, NY. His research interests are in distributed systems, network economics, game theory, and stochastic processes.

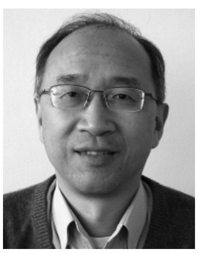

Dah Ming Chiu (F'08) received the B.Sc. degree in electrical engineering from Imperial College London, London, U.K., in 1975, and the Ph.D. degree in applied mathematics from Harvard University, Cambridge, MA, in 1980.

After 20 years in industry (Bell Labs, DEC, and Sun), he is now a Professor with the Department of Information Engineering, Chinese University of Hong Kong. His research interests include network economics and architecture, P2P networking, and wireless networks.

Prof. Chiu is an IEEE Fellow recognized for his "contribution to distributed resource allocation algorithms in computer networks." He serves as an Associate Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING and is regularly invited to serve on the Technical Program Committee of ICNP, INFOCOM, CoNext, IWQoS, and recently SIGCOMM as well.

John C. S. Lui (F'10) was born in Hong Kong and received the Ph.D. degree in computer science from the University of California, Los Angeles, in 1992.

He is currently the Chairman of the Department of Computer Science and Engineering, Chinese University of Hong Kong (CUHK). He worked in the IBM San Jose/Almaden, CA, Lab after his graduation and later joined CUHK. His current research interests are in theoretic/applied topics in data networks, distributed multimedia systems, network security, and mathematical optimization and performance evaluation theory.

Dr. Lui received various departmental teaching awards and the CUHK Vice-Chancellor's Exemplary Teaching Award in 2001. He is a co-recipient of the Best Paper Award in the IFIP WG 7.3 Performance 2005 and the IEEE/IFIP Network Operations and Management (NOMS) Conference. John served as the TPC Co-Chair in the ACM SIGMETRICS 2005 and General Co-Chair of the International Conference on Network Protocols (ICNP) 2007 and is currently the Vice President of ACM SIGMETRICS.

Vishal Misra (S'98–M'99) received the B.Tech. degree from the Indian Institute of Technology, Bombay, India, in 1992, and the Ph.D. degree from the University of Massachusetts, Amherst, in 2000, both in electrical engineering.

He is an Associate Professor of computer science with Columbia University, New York, NY. His research emphasis is on mathematical modeling of computer systems, bridging the gap between practice and analysis. His recent work includes the areas of Internet economics, peer-to-peer networks, and efficient scheduling policies.

Dr. Misra has received an NSF CAREER Award, a DoE CAREER Award, and an IBM Faculty Award. He has served as the Guest Editor for the *Journal of Performance Evaluation* and Technical Program Committee Chair and General Chair of the ACM SIGMETRICS conference. He has participated as a member of program committees for conferences such as the IEEE INFOCM, ACM SIGMETRICS, ACM SIGCOMM, IFIP Performance, and IEEE ICNP. He serves on the Board of Directors of ACM SIGMETRICS.

Dan Rubenstein (M'00) received the B.S. degree in mathematics from the Massachusetts Institute of Technology, Cambridge, in 1992; the M.A. degree in math from the University of California, Los Angeles, in 1994; and the Ph.D. degree in computer science from the University of Massachusetts, Amherst, in 2000.

He is now an Associate Professor of electrical engineering and computer science with Columbia University, New York, NY. He is generally interested in the design and performance of networked systems and also has interests in the security and integrity of such systems as well as the general design of distributed communication algorithms.

Dr. Rubenstein has received a U.S. National Science Foundation CAREER Award, an IBM Faculty Award, an IEEE ICNP Best Paper Award, and (as a student) an ACM SIGMETRICS Best Student Paper Award.