On the Implementation of Path-Based Dynamic Pricing in Edge-Directed Routing

Jumpei Urakawa*, Cristina Basescu†, Kohei Sugiyama*, Christos Pappas†,
Akira Yamada*, Ayumu Kubota*, Adrian Perrig†
*KDDI R&D Laboratories, Inc.
†ETH Zurich

***Abstract*—Future Internet proposals have employed edge-directed routing to realize the benefits of path choice by the sources (e.g., end users). However, economic issues hamper the adoption by ISPs: 1) ISPs' costs increase when sources choose paths that are not economically optimal for ISPs, and 2) ISPs have to overprovision their links aggressively since traffic engineering is shifted to the users and congestion is more likely to occur.**

We implement a path-based dynamic pricing scheme that addresses these challenges. ISPs can dynamically adjust the prices of paths in order to compensate for potential losses incurred by users' choices and to incentivize users to switch paths in case of congestion. We describe our implementation in the context of future Internet architectures and demonstrate a mutually beneficial situation for ISPs and users.

I. Introduction

Edge-directed routing empowers a source to specify a complete or a partial path to a destination. This routing technique comes with multiple benefits: it improves availability and performance through path diversity; it increases competition in the ISP market [1]–[5]; and it enhances security by providing explicit path control [2] [3].

Despite the apparent benefits, edge-directed routing has not been deployed in the Internet, at least partly due to economic issues. Edge-directed routing can raise a conflict of interests between ISPs and sources under today's pricing model: an ISP may have to forward packets over a more expensive provider than the ISP itself would have chosen. Consider the example in Figure 1. $Source_1$ chooses $PATH_1$, which traverses the expensive inter-domain link $AS_2 \rightarrow AS_4$, but AS_2 would prefer the $AS_2 \rightarrow AS_3$ link. Hence, AS_2's cost increases, whereas $Source_1$ continues to pay a flat-rate fee to its provider AS_1. Recent work suggests that a higher flat-rate fee does not always suffice to compensate ISPs for uneconomical user choices [6]. In addition, users prefer the simplicity of flat-rate pricing over other pricing schemes that could perhaps compensate the extra costs for the ISPs [7]. Flat-rate pricing allows users to precisely estimate their broadband expenses compared to other pricing schemes (e.g., volume-based pricing or time-based pricing).

Another issue is that ISPs have to overprovision their links more aggressively to avoid congestion. Since sources have path choice, traffic may converge and congest a downstream link, decreasing the throughput of the flows that traverse it. For example, if many traffic sources choose paths (e.g., $PATH_1$ and $PATH_2$) that traverse the same link ($AS_4 \rightarrow AS_6$), then

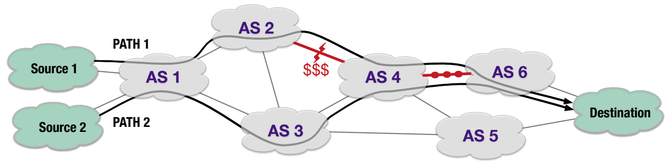

Fig. 1. Edge-directed routing can raise the costs of ISPs through link congestion, and transit of expensive links when cheaper ones are available.

congestion is likely to occur. The root cause of this problem is that traffic engineering is shifted from the ISPs to the end users, but the state of the network remains opaque to the users. AS_4 is in a better position than the end users to handle congestion at the link with AS_6.

Path-based pricing and dynamic pricing may hold the answer to these problems and unleash the potential of edge-directed routing. Agyapong et al. [6] argue that path-based pricing is the only economically viable option for ISPs deploying edge-directed routing. Dynamic pricing is a handy mechanism to address convergence-induced traffic congestion. Namely, ISPs can increase the price of paths that contain congested links and thus incentivize users to switch to cheaper alternatives [8]. Dynamic pricing is known to be an effective method to control users' demands or actions [7], [9]–[11].

Although pricing for edge-directed routing has received considerable research, the proposals so far can be categorized into two disconnected areas: architectures that provide mechanisms and protocols to realize edge-directed routing [1]–[3], [5] and studies that focus on the economic properties of edge-directed routing networks [12]–[15]. The literature lacks an implementation of a pricing scheme for edge-directed routing that bridges this gap.

Furthermore, lacking an implementation, it seems that a pricing scheme in edge-directed routing would either benefit ISPs to the detriment of users or vice versa: either ISPs' revenue increases because the users' cost increases, or the users' cost does not increase and ISPs have a profit loss. *Since one of the parties has to pay an extra cost, is it possible to have an edge-directed routing deployment with a mutually beneficial situation for both ISPs and users?*

Our goal is to provide the first pricing implementation that addresses this seemingly intractable conflict. At the core of our approach is the following intuition: when multiple paths are available, dynamic pricing achieves a better, more balanced network utilization. Thus, users receive overall a higher bandwidth per price unit compared to the bandwidth achieved by today's static pricing schemes, where the price per bandwidth is fixed for a longer term. Congested paths see a price increase, incentivizing some users to switch paths. Thus, higher-priced paths likely turn congestion-free after the price change, and, consequently, become the choice of users that are willing to pay for a better quality of service. ISPs increase their profits because of this additional user payment and because they do not have to overprovision their links.

Specifically, congestion typically occurs at peak times, when users launch bandwidth-hungry applications such as video on demand. These applications usually consist of long-lived flows that are responsible the majority of Internet traffic.[1] Edge-directed routing distributes these high-bandwidth flows over multiple paths, thus avoiding congestion. Since these flows are long-lived, they are also well-suited for a path-based dynamic pricing scheme, where flows switch to cheaper or higher-quality paths, as users prefer. As a result, ISPs reduce their costs by not having to overprovision for peak times, while users maximize their utility by increasing their throughput.

Contributions. We implement a pricing scheme for edge-directed routing that is efficient, considering a selfish behavior from users and ISPs. We describe secure price dissemination and accounting mechanisms in the context of future Internet architectures. We evaluate our proposal under real-world Internet traffic and demonstrate a mutually beneficial situation for both ISPs and users.

II. Problem Description

We first describe the network model and the notation we use. Then, we present our goals.

A. Network Model

The ISP-level topology is an undirected network denoted by $G = (V, E)$, where V is the set of ISPs and E is the set of links. We assume edge-directed routing in the network, where each user $i \in \{1, \ldots, N\}$ generates a traffic flow from a source ISP $s(i) \in V$ to a destination ISP $d(i) \in V$, and selects path $PATH_i$ consisting of links traversed by the flow. Each ISP simply forwards traffic to the next ISP according to the selected path. We also assume that ISPs disclose only price information of their links to users. We divide time into equally sized slots; $t \in \{1, \ldots, T\}$ refers to a time slot.

B. Goals

Our goal is to increase the total utility of users, as well as the total utility of ISPs in the entire network.

Each user $i \in N$ has an increasing, strictly concave, twice differentiable utility function $F_i(\cdot)$, defined over the user's bandwidth $x_i^{(t)}$ at time $t \in T$. We define the total utility of all users in the network as follows:

$$U^{(t)} = \sum_{i \in N} (F_i \left(x_i^{(t)} \right) - x_i^{(t)} p_i^{(t)}), \tag{1}$$

[1]Trammell and Schatzmann [16] have shown that flows with a duration longer than 256ms carry 91.66% of packets.

where $p_i^{(t)}$ is the total path price per bandwidth unit for user i at time t.

Utility functions defined w.r.t. bandwidth are known to be increasing functions [17], [18]. The total utility of users is thus increased if they obtain more bandwidth x. However, the total utility of users, U, could become negative when the users' payment xp exceeds the utility of bandwidth $F(x)$, but we assume users can limit their own sending rate x and lower their expenses. Our scheme aims to increase U by making it economically-viable for ISPs to reduce the price per bandwidth unit, p.

We define the total utility of all ISPs in the network as the aggregate income from selling bandwidth units:

$$R = \sum_{t \in T} \sum_{i \in N} x_i^{(t)} p_i^{(t)}. \tag{2}$$

Since our goal is to increase this total utility, not each ISP's utility, we do not include costs between ISPs.

C. Assumptions

We assume that each AS has a public-private key pair. The public key needs to be known to the other ASes, which can be verified using a PKI infrastructure, such as RPKI. Furthermore, we assume that ISPs are mutually distrustful.

In addition, the source and destination ISPs share a symmetric key, configured through SSH or SSL / TLS, for instance.

III. Path-based Dynamic Pricing

A. Design Decisions

We make the following design decisions to construct a pricing scheme for edge directed routing.

- **Beneficiary pays:** Users pay for the path they select. Since users enjoy the benefits of path diversity, they need to pay for the additional cost of premium paths.
- **ISPs influence path selection:** ISPs can influence users' path selection and do not relinquish all control to users. First, users choose only among a set of paths provided by their ISP, instead of computing arbitrary paths. Second, ISPs influence the path price by changing the price per bandwidth unit of their links.
- **Operational confidentiality:** ISPs disclose only link price information to other ISPs and users. ISPs do not disclose any other information (e.g., congestion status or link capacity) that could reveal business secrets or create new attack vectors.

B. Path-based Dynamic Pricing

We use path-based dynamic pricing that charges users for the selected path, and influences users' path selection by dynamically adjusting the link price. Path-based pricing requires a mechanism to disseminate the path price to all potential users, which we address later in this section.

To support ISPs in controlling users' path selection, our scheme allows each ISP to dynamically change prices on its links. ISPs set a link price per unit of bandwidth, and proportionally charge users based on the traffic sent along paths that traverse their links. Intuitively, as the available link capacity diminishes, ISPs try to avoid link congestion and increase the link price. Thus, the paths traversing the link become more expensive, incentivizing users to select or switch to other paths. Several pricing functions that increase as the link utilization increases achieve our goal. We adopt the congestion-based pricing function defined by Ganesh et al. [9]:

$$F_{ConPrice} = \left(\frac{y}{C}\right)^k, \tag{3}$$

where y is the utilized bandwidth of the link capacity, C is a scale parameter which is associated with the link capacity, and $k \geq 1$ is a scaling parameter for exponential price increase, according to the importance of avoiding congestion for the ISP.

We add to this function a minimum price per bandwidth unit B, to ensure link prices do not drop to 0 when links are not congested. The price of a link becomes:

$$P_l = F_{ConPrice} + B. \tag{4}$$

The resulting path price P_{PATH}, paid by users, is the sum of the prices of all links l on the path:

$$P_{PATH} = \sum_{l \in PATH} P_l. \tag{5}$$

The efficacy of dynamic pricing depends on the time interval to update link prices. As Yuksel et al. show [19], a long update interval hinders ISPs to disperse flows if a link experiences a sudden increase in traffic. Instead, a short interval (1-3 seconds) is desirable to achieve high level of congestion control. However, achieving such a fine-grained price update is considered challenging [7] [11]. When the price propagation delay is higher than the price update interval, it can make users choose paths based on outdated path prices. As a result, the users' traffic may be rejected or it may be charged at a different rate than expected. To solve the issue and achieve a short update interval, we design the effective price dissemination and accounting mechanisms using packet-carried state, as we describe next.

C. Price Dissemination and Accounting

Path-based pricing requires dissemination of path prices to users. Edge-directed routing architectures typically build paths and offer several path choices to users, rather than allowing users to freely specify any sequence of hops [1], [2], [4]. Thus, as a first step, we leverage the path dissemination mechanism to bootstrap price dissemination. However, path dissemination may use a time scale that is too coarse-grained for dynamic pricing. Our solution is *in-path update of prices*: while the source sends traffic along a path, the packets collect new pricing information from the ASes on the path. Because users base their budget on prices, and ASes bill their customers according to prices, prices need to be authenticated by the issuing AS.

We describe price dissemination and accounting in the context of SCION [2], a future Internet architecture that performs edge-directed routing. SCION's key ideas for scalable path construction are to hierarchically group ASes into ISolation Domains (ISDs), and build separate paths in each domain. Then, the source stitches a half-path in the source ISD with a half-path in the destination ISD to form an end-to-end path. Due to lack of space, we explain our scheme in the context of a single ISD (Figure 2), but the same concepts applies for multiple ISDs.

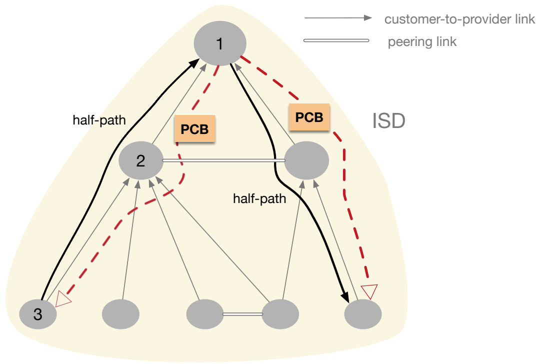

Fig. 2. SCION edge-directed routing using half-paths. Path construction uses path construction beacons (PCBs).

Bootstrapping of price dissemination. Tier-1 ASes in each ISD initiate path construction through Path Construction Beacons (PCBs), propagated down the tree, as depicted in Figure 2. Upon receiving a PCB, AS_i adds its routing information to the PCB (RT_{AS_i}), and authenticates the routing information and the rest of the PCB. Users obtain the routing information and the authenticators from the PCB and insert them in the packet header, to route the packet along the path described by the PCB. The authenticators enable each AS to check whether a packet is allowed to follow the path. The authenticators do not need to be checked by other entities, therefore they are computed by a Message Authentication Code (MAC) with secret key K_i known only to the generating AS_i (Equation 6).

$$Auth_{AS_i} = RT_{AS_i} || MAC_{K_i}\{RT_{AS_i} || Auth_{AS_{i-1}}\}. \tag{6}$$

We leverage PCBs to disseminate the dynamic price of each AS. In contrast to the previously described authenticators, verified only by their originator, prices need to be verified by users, as well as other ASes. We first describe a straw man approach. Each provider AS lists a price and signs it, before appending it to RT_{AS_i}. For example, in Figure 2, AS_1 initiates a PCB, and appends a signature on the dynamic price of the link $AS_2 \rightarrow AS_1$. Then, AS_2 adds to the PCB its routing information and a signature on the dynamic price of the link $AS_3 \rightarrow AS_2$. When a path is used, each AS computes the price to charge its customer summing the prices of all links upstream, after checking their signatures. This is the payment to the AS's provider. For instance, to find the price of the path $AS_3 \rightarrow AS_2 \rightarrow AS_1$, AS_3 sums the prices of the links $AS_2 \rightarrow AS_1$, and $AS_3 \rightarrow AS_2$, and checks their signatures.

Computing and checking signatures is computationally expensive.[2] Since signatures are generated with every incoming beacon, we must bound the number of propagated beacons. SCION builds most paths that exist in today's Internet when beacons are created every 15 seconds, and each AS forwards at most 5 beacons. However, price computation during data-traffic forwarding must use symmetric cryptographic operations for price authenticity, in order to support high forwarding speeds. We optimize data packet forwarding by letting routers *compute the path price during beacon propagation*. In Figure 2, AS_1 simply inserts the price and the signature, as before. However, when AS_2 receives the beacon, it first checks the signature of the price $AS_2 \rightarrow AS_1$, and puts a MAC using its own key K_2. Then it adds to the price of $AS_2 \rightarrow AS_1$ the price of the link $AS_3 \rightarrow AS_2$, and signs the total price. AS_3 verifies the signature of AS_2 for the price of the path $AS_3 \rightarrow AS_2 \rightarrow AS_1$, and puts a MAC with its own key K_3. An customer of AS_3 that uses the path simply pays the path price plus a service fee, communicated by its service provider AS_3. During packet forwarding, each AS only needs to check the MAC it generated in order to correctly pay its AS provider. The procedure is illustrated in Equation 7.

$$\begin{aligned} Auth_{AS_i} = {} & RT_{AS_i}||Price_{AS_i \rightarrow Dst}||Price_{AS_{i-1} \rightarrow Dst} \\ & ||Sign_{K_i^{-1}}\{Price_{AS_{i-1} \rightarrow Dst}\}|| \\ & ||MAC_{K_i}\{RT_{AS_i}||Price_{AS_i \rightarrow Dst}||Auth_{AS_{i+1}}\}, \end{aligned} \tag{7}$$

where AS_{i-1} and AS_{i+1} are the customer and the provider of AS_i respectively.

In-path price update. Even though ASes forward SCION beacons every few seconds, the time required to update a PCB for *a particular path* may be in the order of minutes. The reason is the large number of existing beacons and the limit of 5 beacons propagated per link.

To achieve price update at the granularity of a second, packets traversing a particular path collect updated pricing information from the ASes on the path. Using the same example in Figure 2, when a packet travels towards AS_3, AS_1 adds to the packet header a new price for the link $AS_2 \rightarrow AS_1$, together with a validity period t_1; likewise, AS_2 proceeds in the same way for the link $AS_3 \rightarrow AS_2$. Each AS needs to authenticate the prices for the source. The destination feeds back the new pricing information to the source.

To authenticate these prices, we make use of the DRKey protocol [21], where a source shares secret keys with each AS on the path. The main benefit of DRKey is to allow each AS_i to rapidly derive on-the-fly a key shared with each source, K_{Si}, instead of storing these keys as per-flow state, which would require excessive storage on core routers.

Each AS uses the derived key to compute a MAC over its price. The source can then sum the prices and check their authenticity, as well as check the validity of prices through the timestamps t_i. The AS also authenticates with its key the payment it should receive from its customer AS_{i-1}, and the price's validity (Equation 8). The update inserted by AS_i in the packet header is given by Equation 9.

$$\begin{aligned} P_{own-i} = {} & Price_{AS_{i-1} \rightarrow AS_i}||t_i|| \\ & ||MAC_{K_i}\{Price_{AS_{i-1} \rightarrow AS_i}||t_i\}, \end{aligned} \tag{8}$$

$$Update_i = P_{own-i}||MAC_{K_{Si}}\{Price_{AS_{i-1} \rightarrow AS_i}||t_i\}. \tag{9}$$

After price updates, the source can trivially compute the correct path price, however, ASes on the path cannot compute the sum to pay their provider, because they cannot check the authenticity of all link prices. For this reason, the source computes for each AS_i two values: P_{rcv-i} is the payment AS_i receives from its customer on the path, and P_{own-i} is the sum AS_i keeps for itself from the payment. The source authenticates these values for AS_i using a shared symmetric key K_{Si}.

$$\begin{aligned} Payment_i = {} & P_{rcv-i}||P_{own-i}||t_i|| \\ & ||MAC_{K_{Si}}\{P_{rcv-i}||P_{own-i}||t_i\}. \end{aligned} \tag{10}$$

To illustrate the procedure, in Figure 2, an end-host in AS_3 receives a price update for the path $AS_3 \rightarrow AS_2 \rightarrow AS_1$. The end-host computes $P_{rcv-3} = Price_{AS_3 \rightarrow AS_2 \rightarrow AS_1} + fee$, $P_{own-3} = fee$, $P_{rcv-2} = Price_{AS_3 \rightarrow AS_2 \rightarrow AS_1}$, $P_{own-2} = Price_{AS_3 \rightarrow AS_2}$, $P_{rcv-1} = Price_{AS_2 \rightarrow AS_1}$, $P_{own-1} = Price_{AS_2 \rightarrow AS_1}$, where *fee* is a service fee, communicated by end-host's service provider. Afterwards, the end-host authenticates these values as described by Equation 10.

To settle payments, each AS sums per link the authentic P_{own} payments received on that link. Packets with invalid P_{own} fields are dropped. The storage needed at the AS to perform accounting is very small: only one value per neighbor. The AS charges the total sum from each neighbor periodically (e.g., at the end of the day or once a month).

IV. Evaluation

A. Simulation Setup

To quantify the effect of dynamic pricing in edge-directed routing, we compare dynamic pricing with static pricing, both per bandwidth unit, and evaluate the total utility of users and ISPs (Section II-B).

User Behavior. Users send a given amount of traffic along a selected path – we refer to it as a flow – starting at a given point in time. We assume that users choose the cheapest path they know to a destination, and switch the path when they consider a flow to be long-lived. For short-lived flows, path switching overhead, such as connection establishment and TCP slow-start, can outweigh the benefits of switching to a cheaper path. For our simulation, users switch to a cheaper path if the estimated completion time of the flow exceeds 2 seconds.

Network Topology and Traffic Generation. We generate a network topology based on a CAIDA dataset [22]: we select

[2] ed25519 [20] provides fast signing and verification (about 71000 verifications, and 109000 signatures generated per second), yet this is too slow for core routers, forwarding billions of packets per second.

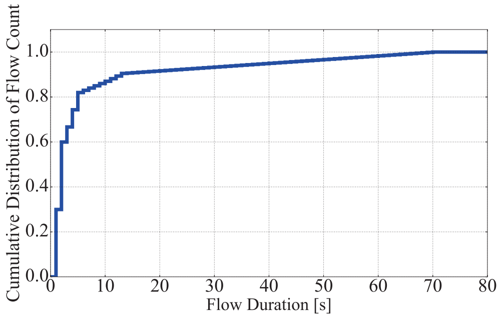

Fig. 3. Distribution of input flows' duration.

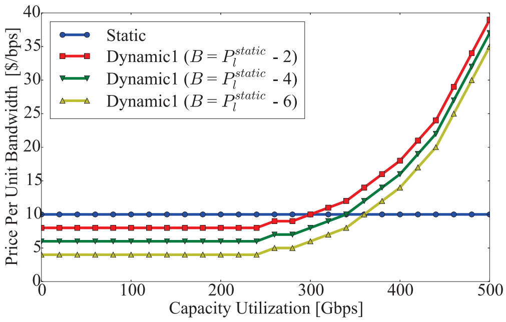

Fig. 4. Comparison between static and dynamic pricing functions. The threshold congestion for dynamic pricing is 250 Gbps.

615 ISPs managed by JPNIC [23], and assume that their interconnection links have a capacity of 500 Gbps.[3]

Furthermore, we generate traffic patterns according to statistics from a recent survey [16]: there are 10,000 to 20,000 concurrent flows per /16 address block (i.e., 65534 hosts) during a 5-second window. Since JPNIC manages 93,059,006 hosts [23], there are 14,200,100 - 28,400,220 flows per 5 seconds interval. We choose a value in this interval, 21,300,00 flows per 5 seconds, or equivalently 4,260,000 flows per second. We generate the flows according to a Poisson distribution, and set the flow durations according to the same survey as Figure 3. [4] Sources and destinations are chosen uniformly at random.

Link Prices. We set a random static link price from 10 to 15 units (denoted by P_l^{static}). For the dynamic link price, we set the parameter $k = 5$ (Equation 3). We set a minimum price B smaller than P_l^{static}, and let the price dynamically increase when the link utilization exceeds 50%, i.e., exceeds 250 Gbps. We denote $\alpha \leftarrow P_l^{static} - B$, and vary α with 2, 4, and 6 price units.

The behavior of the link price per bandwidth unit is depicted in Figure 4. The figure shows that, once the link utilization

[3] A link represents an interconnection between two ISPs and does not correspond to physical connectivity; two ISPs can interconnect at multiple places.

[4] Since the survey does not show detailed distribution of flow durations, we generate approximation based on described data.

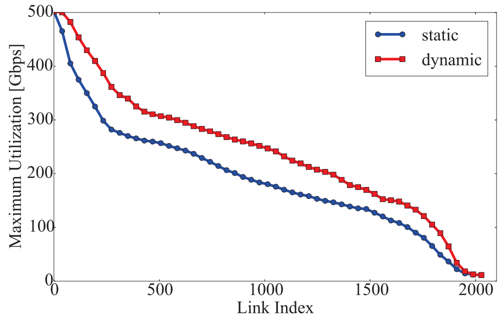

Fig. 5. Maximum utilization of links during the observation period for static and dynamic pricing.

ratio exceeds 50%, the ISPs' dynamic link price exceeds the static link price, so as to avoid congestion.

B. Results

We compare the network utilization under static and dynamic pricing. We also analyze the bandwidth and the price per unit bandwidth, which constitute users' utility, and the revenue, which constitute ISPs' utility. In addition, we show the frequency of path switching related to network stability.

Link Utilization. Figure 5 shows that under dynamic pricing, network links are consistently higher utilized, as sources prefer paths traversing lesser-utilized links, because they are cheaper than congested paths. Our results show that dynamic pricing utilizes the network resources more efficiently than static pricing, increasing the bandwidth of flows by as much as 15% compared to static pricing.

Bandwidth Unit Price and ISPs' Utility. Table I compares static and dynamic pricing w.r.t. the average bandwidth flows obtain, the price per bandwidth unit, and the ISPs' revenue. We set the static link price $P_l^{static} = 35$, and set the minimum price B for the dynamic pricing as before, smaller with 2, 4, and 6 units than P_l^{static}. The results show that the pricing function determines whether the ISPs or the users obtain a higher utility. For instance, in the case $\alpha = 6$, users obtain a smaller price per bandwidth unit than in the cases $\alpha = 2$, $\alpha = 4$, and static pricing, but the ISPs' revenue is the smallest among all cases. Conversely, when $\alpha = 2$, the ISPs obtain the highest revenue, but the users' price per bandwidth unit is the highest. However, there are dynamic pricing functions, such as $\alpha = 4$, which at the same time increase the ISPs' revenue and decrease the users' price per bandwidth unit, compared to static pricing. Therefore, dynamic pricing can make edge-directed routing more profitable.

Frequency of Path Switching. Figure 6 shows the distribution of the number of path switches for dynamic pricing. As shown in this figure, approximately 90% of users do not change the path because their flows have a short duration, and the switching overhead outweighs the benefits of the cheaper path. Although a few users often switch the path due to limitation of available paths to reach the destinations, as shown in this

TABLE I
AVERAGE FLOW BANDWIDTH, PRICE PER BANDWIDTH UNIT, AND ISPS' REVENUE. PU DENOTES PRICE UNIT.

	Static P_l^{static}	Dynamic ($B = P_l^{static} - \alpha$) $\alpha = 2$	$\alpha = 4$	$\alpha = 6$
Avg.Bw. [Mbps]	8.2	9.1	9.1	9.0
Bw. unit price [PU/bps]	35	40	34	28
Revenue [Mil.PU]	470,272	607,486	511,232	420,047

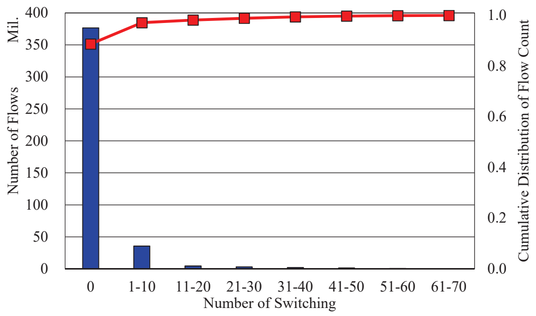

Fig. 6. Distribution of the number of switching for dynamic pricing.

result, the flows are a small part of entire flows (a few percentage) and have little impact for entire network.

V. Discussion

Future capacity expansion costs of ISPs. In edge-directed routing, since it is difficult for ISPs to estimate how many users select their links, ISPs are forced over-provisioning for their capacity. However, if edge-directed routing implements dynamic pricing, ISPs can achieve higher link utilization, because traffic finds alternate paths without causing congestion. We take the example of links whose maximum utilization is 50%, which are overprovisioned links. The number of such links is 1478 (72%) in static pricing and 1057 (52%) in dynamic pricing (Figure 5). This means that dynamic pricing could diminish the wasteful costs of ISPs on the entire network by 20%. In addition, ISPs will not need to spend the excessive costs for capacity expansion because they can control users' demands by increasing the unit price.

Network Stability. Since dynamic pricing changes user's demand for each path, the dynamics possibly causes excessive path switching. As shown in Section IV, such flows account for only a small part of entire network under real-world Internet environment. However, emergence of massive users who do not take into account own switching overhead or who try to attack this pricing system perhaps causes high-frequency switching or fluctuation of link price. To mitigate such effects and keep network condition more stable, ISPs can incorporate feedback control technique into proposed mechanism for price determination (Equation 4). For example, ISPs set the link price at time t using price in previous time slot as follows:

$$P_l^{(t)} = (1-s)P_l^{(t-1)} + s(F_{ConPrice} + B), \quad (11)$$

where $s \in [0,1]$ is the value of price sensitivity. ISPs also can reduce frequency of price change by extending the time interval to update link prices. Since these approaches include a trade-off relationship between immediacy and stability, ISPs which prefer to handle burst traffic immediately make price sensitivity higher (i.e., increase s) or interval shorter, conversely, they can set lower value for sensitivity and longer interval if they prefer to provide more stable and user-friendly price.

VI. Related Work

We briefly describe related work in the area of pricing schemes for edge-directed routing and dynamic pricing in general.

Pricing for edge-directed routing. There is limited work in the area of pricing for edge-directed routing. Agyapong and Sirbu [6] show that an increase of traffic generated by end users suffices to compensate for the additional costs incurred by uneconomical path choices by users. Furthermore, if users' choices result in path oscillation, then path-based pricing is the only economically viable solution for ISPs. However, the paper does not study the effects of traffic convergence on users' throughput and ISPs' revenues.

Laskowski et al. [8] argue that selfish user's behavior can be dealt with dynamic prices. This basic idea is similar to ours. However, as these authors also described in the paper, this study is lacking a deployment path such as who decides which path's price, how to decide price (i.e., Based on what price is decided), how to achieve dissemination and accounting and so on. In addition, evaluation with a realistic network topology and actual traffic patterns is also missing.

Dynamic pricing. Dynamic pricing schemes have received considerable attention from the research community. Ganesh et al. [9], [24] use congestion-based pricing to achieve a fair and efficient bandwidth utilization. The authors show that the users' price estimations converge to the actual cost, guaranteeing stability for the network.

Paschalidis and Tsitsikilis [10] address revenue and welfare maximization under congestion-based dynamic pricing. They show that *time-of-day pricing*, which varies prices according to users' demand during the day, is close to the optimal solution with respect to revenue and welfare maximization.

Sen et al. [7], [11], [25], [26] argue that *day-ahead pricing*, which determines prices based on users' previous day's traffic patterns, is more user-friendly than congestion-based pricing due to longer price update intervals. Loiseau et al. [27] proposed *raffle-based pricing* inspired by lottery reward mechanism. This scheme gives users a reward for reducing their usage of the shared resource. Since these pricing schemes decide price before (i.e. 1 day ago for day-ahead pricing) or after (i.e. after consuming a certain amount of data for raffle-based pricing) actual resource usage, implementation of price dissemination and accounting is relatively easy. However, these pricing scheme could not deal with sudden and unexpected increases of traffic as we described in Section V.

VII. Conclusion

Edge-directed routing has been known to provide benefits for the end users, but it has not been deployed due to economic and performance issues. This paper has presented the implementation of a path-based dynamic pricing scheme that addresses these issues. To protect ISPs' economic interests when users choose paths that ISPs do not prefer, the extra cost is shifted to the end users, who are the beneficiaries of path selection. Furthermore, to avoid congestion due to traffic convergence on certain paths, ISPs can influence the price of a path and incentivize sources to choose another one. With our implementation, we have demonstrated that this pricing scheme is mutually beneficial for ISPs and users: ISPs increase their revenue and users enjoy a higher bandwidth at a lower price per bandwidth unit.

References

[1] X. Yang, D. Clark, and A. W. Berger, "Nira: A new inter-domain routing architecture," *IEEE/ACM Transactions on Networking*, 2007.

[2] X. Zhang, H.-C. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. G. Andersen, "Scion: Scalability, control, and isolation on next-generation networks," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2011.

[3] C. Basescu, R. M. Reischuk, P. Szalachowski, A. Perrig, Y. Zhang, H.-C. Hsiao, A. Kubota, and J. Urakawa, "Sibra:scalable internet bandwidth reservation architecture," in *Proceedings of Network and Distributed System Security (NDSS)*, 2016.

[4] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica, "Pathlet routing," in *Proceedings of the ACM SIGCOMM Conference on Data Communication*, 2009.

[5] R. Raghavan and A. C. Snoeren, "A system for authenticated policy-compliant routing," in *Proceedings of the ACM SIGCOMM Conference on Data Communication*, 2004.

[6] P. K. Agyapong and M. Sirbu, "The economic implications of edge-directed routing: A network operator's perspective," in *Proceedings of the Fourth International Conference on Evolving Internet (INTERNET)*, 2012.

[7] S. Sen, C. Joe-Wong, S. Ha, and M. Chiang, "A survey of smart data pricing: Past proposals, current plans, and future trends," *ACM Computing Surveys (CSUR)*, 2013.

[8] P. Laskowski, B. Johnson, and J. Chuang, "User-directed routing: From theory, towards practice," in *Proceedings of the 11th workshop on the Economics of Networks, Systems and Computation (NetEcon)*. ACM, 2008.

[9] A. Ganesh, K. Laevens, and R. Steinberg, "Congestion pricing and user adaptation," in *Proceedings of IEEE Computer and Communications Societies (INFOCOM)*, 2001.

[10] I. C. Paschalidis and J. N. Tsitsiklis, "Congestion-dependent pricing of network services," *IEEE/ACM Trans. Netw.*, 2000.

[11] S. Ha, S. Sen, C. Joe-Wong, Y. Im, and M. Chiang, "Tube: Time-dependent pricing for mobile data," 2012.

[12] A. Orda, R. Rom, and N. Shimkin, "Competitive routing in multiuser communication networks," *IEEE/ACM Trans. Netw.*, vol. 1, no. 5, pp. 510–521, Oct. 1993. [Online]. Available: http://dx.doi.org/10.1109/90.251910

[13] K. Park, M. Sitharam, and S. Chen, "Quality of service provision in noncooperative networks with diverse user requirements," *Decision Support Systems*, vol. 28, pp. 101 – 122, 2000. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167923699000780

[14] T. Roughgarden and E. Tardos, "How bad is selfish routing?" *J. ACM*, vol. 49, no. 2, pp. 236–259, Mar. 2002. [Online]. Available: http://doi.acm.org/10.1145/506147.506153

[15] L. Qiu, Y. R. Yang, Y. Zhang, and S. Shenker, "On selfish routing in internet-like environments," in *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '03. New York, NY, USA: ACM, 2003, pp. 151–162. [Online]. Available: http://doi.acm.org/10.1145/863955.863974

[16] B. Trammell and D. Schatzmann, "On flow concurrency in the internet and its implications for capacity sharing," in *Proceedings of the ACM Workshop on Capacity Sharing (CSWS)*, 2012.

[17] F. Paganini and E. Mallada, "A unified approach to congestion control and node-based multipath routing," *IEEE/ACM Transactions on Networking*, vol. 17, no. 5, pp. 1413–1426, Oct. 2009.

[18] S. H. Low and D. E. Lapsley, "Optimization flow control, I: Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, Dec. 1999.

[19] M. Yuksel and S. Kalyanaraman, "Effect of pricing intervals on congestion-sensitivity of network prices," *Telecommunication Systems*, 2005.

[20] D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang, "High-speed high-security signatures," in *Proceedings of Cryptographic Hardware and Embedded Systems (CHES)*, 2011.

[21] T. H.-J. Kim, C. Basescu, L. Jia, S. B. Lee, Y.-C. Hu, and A. Perrig, "Lightweight source authentication and path validation," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, 2014.

[22] "Center for Applied Internet Data Analysis (CAIDA)," http://www.caida.org/home/, 2015.

[23] "Japan Network Information Center (JPNIC)," https://www.nic.ad.jp/en/, 2015.

[24] A. Ganesh, K. Laevens, and S. Richard, "Dynamics of congestion pricing," *Microsoft Research Technical Report MSR-TR-2000-70*, 2000.

[25] S. Sen, C. Joe-Wong, S. Ha, J. Bawa, and M. Chiang, "When the price is right: Enabling time-dependent pricing of broadband data," in *Proceedings of ACM Special Interest Group on Computer-Human Interaction (SIGCHI)*, 2013.

[26] C. Joe-Wong, S. Ha, and M. Chiang, "Time-dependent broadband pricing: Feasibility and benefits," in *Proceedings of IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2011.

[27] P. Loiseau, G. Schwartz, J. Musacchio, and S. Amin, "Incentive schemes for internet congestion management: Raffles versus time-of-day pricing," in *Proceedings of the Allerton Conference*, 2011.

[28] P. Hande, M. Chiang, R. Calderbank, and J. Zhang, "Pricing under constraints in access networks: Revenue maximization and congestion management," in *Proceedings of the 29th Conference on Information Communications (INFOCOM)*, 2010.

[29] I. Castro, A. Panda, B. Raghavan, S. Shenker, and S. Gorinsky, "Route bazaar: Automatic interdomain contract negotiation," in *15th Workshop on Hot Topics in Operating Systems (HotOS)*, 2015.