Exit Policy Violations in Multi-hop Overlay Routes: Analysis and Mitigation

Srinivasan Seetharaman and Mostafa Ammar
Networking and Telecommunications Group, College of Computing
Georgia Institute of Technology, Atlanta, Georgia 30332
{srini,ammar}@cc.gatech.edu

***Abstract*— The traffic exchanged between two overlay nodes in different autonomous systems (AS) is always subjected to a series of inter-domain policies. However, overlay routing often manages to get around these policy restrictions by relaying traffic through multiple legitimate segments, in order to achieve its selfish goals (e.g., better latency paths between end-systems). We focus on the violation of a generalized *exit policy*, which specifies the exact next hop AS and the egress inter-domain link for a destination address prefix. We characterize the different types of these exit policy violations and investigate their extent in a Planetlab testbed. It is conceivable that the native ASes will eventually realize the negative impact of the exit violations and adopt stringent strategies to enforce the exit policies, thereby causing deterioration in overlay performance. In this context, based on our findings from a previous study[1], we develop a pricing-based strategy that an overlay service provider can use to obtain permits from a near-optimal set of native ASes, in an effort to regain its routing advantage within a fixed budget. Further, we illustrate the use of this approach on our case study overlay network.**

I. Introduction

Overlay networks have recently attained popularity as a way to deploy functionality that would generally require substantial modification at the IP layer. By forming a virtual network on top of the physical network, the overlay nodes collaborate to offer specialized services like multicast[2] and QoS[3]. The route traversed by each overlay packet is configured to achieve a specific end-to-end performance objective. In doing so, it is possible that the overlay packet is *relayed* through one or more intermediate overlay nodes. We refer to the route between each overlay neighbor as an *overlay link* and the end-to-end route between two overlay nodes as an *overlay path*[1].

Based on the location of the individual overlay nodes, the overlay path can span multiple native autonomous systems (AS), each with its own inter-domain policy restrictions. In our previous work[1], we investigated the extent of inter-domain transit policy violations caused by overlay routing. Specifically, we focused on the overlay paths that violate the *valley-free* property of certain ASes, which states that no AS will act as a transit for traffic originating from its provider or its peer, unless the traffic is destined to its customer[4]. By constructing a case study overlay network over Planetlab[5] and other synthetic topologies, we estimated the extent of transit violations committed when the objective of the overlay network was to optimize the end-to-end latency between its nodes. We observed that nearly 70% of the multi-hop overlay paths achieved better routing performance over the direct native route by adopting routes that violate the valley-free property. It is worth noting that the brunt of these transit violations are experienced by an intermediate AS not involved in any way with the end-to-end communication.

This work was supported in part by NSF grant ANI-0240485.

[1]Our work pertains to infrastructure or service overlays, rather than peer-to-peer networks.

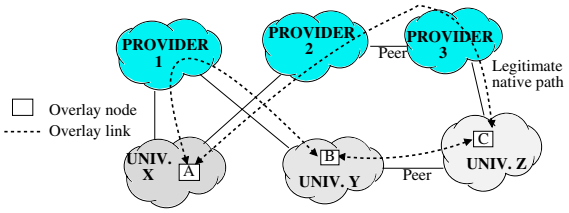

Fig. 1. Exit policy violations caused by overlay routing.

In this paper, we investigate the policy violations caused by overlay routing from a broader perspective, in an effort to cover all possible policy violations. The analysis is based on our observation that any deviation from the legitimate route between two overlay nodes can be traced down to a change in the egress point somewhere upstream. We refer to the broader set of violations as *exit policy violations*; wherein the exit policy of a particular AS specifies i) the exact next hop AS, and ii) the egress inter-domain link for a destination address prefix. These violations can occur at any point of the route.

We are concerned about these violations because they lead to unfavorable economics for an AS and undesirable increase in load across certain inter-domain links. Fig. 1 illustrates these problems in a hypothetical AS-level connectivity graph. In that figure, overlay node A can route data to node C using the overlay path *ABC*, which causes University X to pick Provider 1, rather than the conventional Provider 2. This is a violation of University X's exit policy, though it goes undetected by the native layer. From an economic perspective, we see that University X potentially has to spend more money to route traffic through Provider 1.

We present our classification of the different types of exit policy violations and describe their implications. Further, by analyzing the routes used in a case study overlay network, we provide insights into the frequency and extent of each type. We observed that over 87% of the multi-hop overlay paths in our dataset violate the exit policies of the native ASes.

These exit policies, typically motivated by economic and performance gains, reflect the commercial agreements between an AS and its neighbors[6]. Hence, a violation of these policies by overlay routing can have potentially serious implications, which are objectionable to the ASes. Thus, it is conceivable that the ASes will adopt various strategies to control and manage the overlay traffic, leading to deterioration in user experience[7], [8]. We showed in our previous work a clear tradeoff between the number of ASes enforcing these policies and the penalty incurred by the overlay performance. This motivates the need for a solution that is agreeable to both the native ASes and the overlay network.

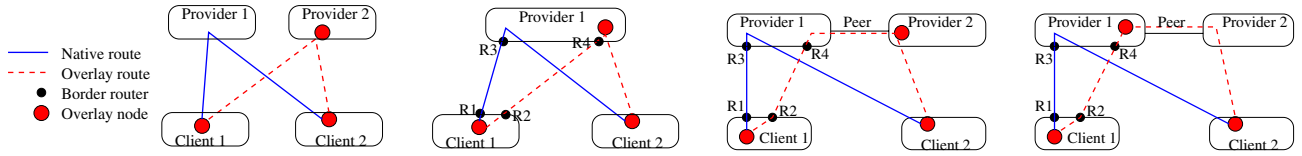

E1. Next hop AS violated

E2. Localpref/MED violated causing exit point change

E3. Next hop AS violation causing exit point violation

E4. Exit point violation causing a next hop AS violation

Fig. 2. Possible exit policy violations. Each of these violations can occur at any point of the multi-hop path, either at a host AS or a non-host AS.

For cases where there exists a common overlay service provider (OSP) operating the overlay network, we propose a pricing-based strategy that the OSP can use to achieve the best possible routing performance within a fixed budget. In particular, the OSP uses this budget to obtain permits from specific ASes to condone the exit policy violation and to allow overlay traffic to use its preferred exit point. However, resolving objections to the undesired exit policy violations does not necessarily legitimize a multi-hop overlay path, as the path may still be ridden with transit policy violations. Therefore, our pricing strategy typically serves to augment the *cost-sharing* approach proposed in [1].

The remainder of the paper is organized as follows. We describe the different types of exit policy violations possible in Section II. We characterize the extent of these policy violations using our case study overlay network and present associated results in Section III. Section IV presents our pricing-based strategy for resolving the conflict between the two layers and arriving at a solution that is mutually agreeable to both. Previous research related to our work is briefly described in Section V. This paper is concluded in Section VI.

II. A Classification of Exit Policy Violations

The Internet is formed by the interconnection of numerous autonomous systems (AS), wherein each AS adopts a certain inter-domain policy to control the route of traffic entering and exiting its network. This leads to the conception of the AS relationships, which reflect the commercial agreements between an AS and its neighbor. In our previous work[1], we studied how overlay routes violate these AS relationships by relaying traffic through intermediate overlay nodes. We noted that all multi-hop overlay paths represent a violation of the valley-free property, unless the intermediate relay is located in a provider network.

In this paper, we study a broader class of inter-domain policy violations, where the exit policy of a particular AS is violated. The exit policy is modeled as *the preferred combination of the next hop AS and the egress inter-domain link, for a particular destination IP prefix*. We define exit policy violation as a deviation from this preferred exit, caused by overlay relaying that hides the actual destination from the native layer. We posit that such deviation causes undesired load and expenses for the native layer. The level of objection to each exit violation may vary with each AS and each overlay path. We focus specifically on multi-hop overlay paths as single-hop overlay paths, which follow the same route as the direct native route, do not violate any native layer policy; under the assumption that the native routing is always policy-compliant.

Fig. 2 illustrates the four basic forms of exit policy violations possible. We describe each as follows:

E1. *Next hop AS violated*: This is caused when the overlay traffic is relayed through an intermediate node located in an AS not along the direct native route between the end nodes. Fig. 2(E1) illustrates a simple scenario where the source overlay node causes Client_1 to pick Provider_2 to indirectly reach a destination in Client_2.

E2. *Ingress or Egress router preference violated*: An exit point violation can happen when an overlay path uses a relay node in a downstream AS that is closer to a different ingress router not used by the direct native route. Clearly, this could be a violation of the Localpref attribute, hot-potato routing or cold-potato routing. In Fig. 2(E2), we can see that the local egress preference (router $R1$) and the neighbor's ingress preference (router $R3$) are violated, without the knowledge of either AS. This has the problem that the load from the overlay traffic is borne by the link $R2-R4$ instead of the designated inter-domain link $R1-R3$.

E3. *Exit point violated because of a next hop AS violation*: This form is similar to the previous scenario, except that the router preference is affected by the alteration of the next hop AS at a downstream provider. In Fig. 2(E3), we see that a change in next hop AS at Provider_1 causes a change in the ingress point from $R3$ to $R4$. Such a change in ingress router is achieved when Provider_1 offers a different MED value for its perceived destination prefix.

E4. *Next hop AS violated because of an exit point violation*: When a downstream AS spans a wide geographical region, it is possible that a change in the ingress point into its domain might cause it to alter its preference of the next hop AS. In Fig. 2(E4), we see that a change in the ingress point from $R3$ to $R4$ causes Provider_1 to transit traffic through Provider_2, rather than sending directly to Client_2.

The above four scenarios capture the different types of exit policy violations. Each violation has serious economic or load repercussions and is undesirable from the perspective of the native service provider.

An interesting artifact of shortest path routing performed at the overlay layer is that *it is sufficient to analyze each 2-hop overlay path*. Consider an overlay network with nodes A, B, C and D. If the shortest path between nodes A and D is $ABCD$, then the shortest path between nodes A and C is ABC. Hence, if the 2-hop overlay path ABC or BCD is violating, then the 3-hop overlay path $ABCD$ will also be violating. Furthermore, the number of violations in a multi-hop path is a summation of the violations observed in its constituent 2-hop overlay paths. This confirms that the 2-hop overlay path scenarios considered in Fig. 2 represent all types of exit violations possible.

The exit violations in each path originate at a single AS where the inter-domain link changes. However, based on the particular type, it may be experienced by other downstream

ASes as well. Unlike valley-free violations, the exit violations are not restricted to only an intermediate host AS[2]. From the classification in the previous subsection, we observe that an exit violation can originate at any one of the three following locations:

- Source AS (which is also a host AS)
- Intermediate host AS
- Intermediate non-host AS

We argue that any overlay path having a valley-free violation (discussed in [1]) necessarily has at least one exit violation at an upstream AS. This can be reasoned by the fact that a valley-free violation occurs at an intermediate host AS that lies outside of the direct native route and such a deviation must originate at an upstream AS (See Fig. 1, for example). Thus, exit violating paths represent the superset of all violating paths. Furthermore, in a particular multi-hop overlay path, the same AS will never experience both exit violations and valley-free violations. This is because valley-free violations happen only at ASes that are not along the legitimate route between the end points, in contrast to exit violations.

III. Characterizing Exit Policy Violations

In this section, we provide insights into the extent of exit policy violations, by means of an experimental case study overlay network constructed over Planetlab[5].

A. Overlay Network Case Study

Our case study testbed was made up of 58 Planetlab nodes that are geographically distributed (based on latitude/longitude) over North America, with only one node per AS. Thus, the total number of overlay paths we inspected for inter-domain policy violations was 3306. We assume complete mesh connectivity of overlay links in our testbed, for all results presented in this paper. This gives a good degree of freedom in determining overlay paths.

For our experiment, we deployed a service overlay network with an objective of offering lowest latency routes between each pair of overlay nodes. We picked the latency metric because it is a crucial factor for most end-system applications. Nevertheless, we posit that overlay routing will lead to policy violations for any choice of routing metric, since the basic requirement for violations to be caused is the presence of multi-hop overlay paths. When we performed shortest path routing with the link metric being latency, we found it beneficial to relay the overlay traffic through one or more intermediate overlay nodes in 56.5% of the overlay paths, though there exists a direct native route between each pair of overlay nodes (represented by the overlay link). Refer to [1] for further details on the characteristics of the overlay paths derived.

B. Measurement Methodology

To estimate the number of exit policy violations observed in the multi-hop overlay paths, we adopt the following steps:

- Measure the latency across each overlay link using ping.
- Obtain the shortest path between each pair of overlay nodes as a sequence of the individual overlay links traversed.
- Determine the route taken by each overlay link using the Scriptroute[9] tool[3]. This was feasible because we had direct access to each of the 58 overlay nodes.
- Map the IP address of each hop in the overlay link to its corresponding AS number, using the publicly available IP-prefix-to-AS mapping generated by the algorithm in [10].
- Combine all this information to obtain the *end-to-end AS path*, defined as the sequence of ASes traversed by the overlay traffic.
- Analyze the sequence of inter-domain links used by the direct native route and the actual overlay path. Any variation in the inter-domain link (the border router and/or the AS number at each end) represents an exit violation. Consequently, we classify them according to the four different categories based on the exact router and AS hops.
- Infer the AS relationships using Gao's algorithm[4] and some simple corrections to reduce inaccuracies[1]. This requires a more complete view of the BGP routes used in the Internet. As suggested in [11], we extracted the BGP tables from 6 RouteViews servers, 14 RIPE RCCs, 30 public routeservers and 1 lookingglass server. Using the AS relationships inferred from the BGP routes, we estimate the number of valley-free violations present in each end-to-end AS path.

[2]We refer to an AS in which an overlay node is located as the *host AS*.

[3]In certain anomalous cases where the rockettrace function did not work, we used the traceroute tool.

TABLE I

Exit violations noticed in the Planetlab dataset, along with the number of paths having valley-free violations.

Type	Originating Location	Exit violations	%	Valley-free violations
E1	Source AS	502	26.9	444
	Intermediate host AS	104	5.56	76
	Intermediate non-host AS	372	19.9	331
E2	Source AS	43	2.30	1
	Intermediate host AS	113	6.05	0
	Intermediate non-host AS	135	7.22	11
E3	Intermediate host AS	26	5.83	19
	Intermediate non-host AS	342	18.3	326
E4	Source host AS	1	0.05	0
Total violating paths		1638	87.7	1208

C. Policy Violations Observed

Table I presents the frequency of exit violations noticed by comparing the shortest path and the direct native path in the original Planetlab data. We note that nearly 87.7% of the multi-hop overlay paths represent an exit violation. Interestingly, not all multi-hop overlay paths represent an exit violation as one might expect. This is because the AS path and the exit routers used are the same for almost 12.2% of the overlay paths, though the exact route traversed by the multi-hop overlay path is indeed different from that of the direct native route. This peculiarity is mainly observed when the intermediate node is located in the Internet2 AS. We also affirm from the data that all 1208 multi-hop overlay paths having valley-free violations also have exit violations.

Most of the exit policy violations we observed in our testbed network was of the type where the next hop AS is changed by overlay routing, viz. $E1$ and $E3$. Note that for violations of type $E3$, we consider the change in the next hop AS as the most reproachable and mark the AS experiencing it as the origin point. Hence, we do not see a source AS originating the violation $E3$. Furthermore, we see more exit violations originating at source ASes (33.2% of total) and at intermediate non-host ASes (51.8% of total) in our dataset.

As mentioned earlier, a violation in a 2-hop overlay path BC is potentially seen in each overlay path AD that overlaps path BC i.e., the same exit violation might be part of two dif-

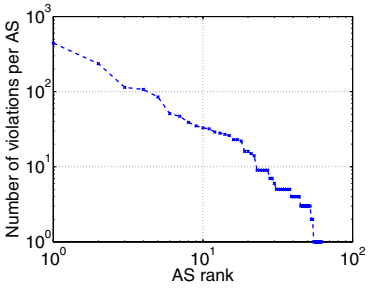

Fig. 3. Log-Log plot of the number of exit violations experienced by each of 62 ASes. This shows clear non-uniformity in the violations experienced.

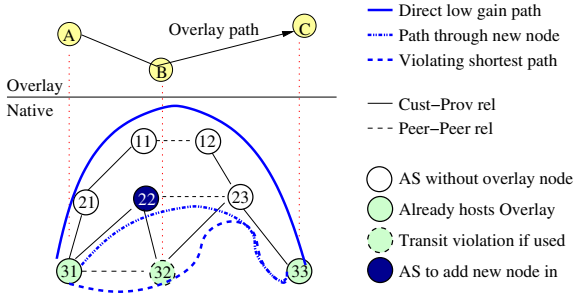

Fig. 4. Cost-sharing approach to mitigate the valley-free transit violations.

ferent paths because the end points are different, although the intermediate segments are the same. Furthermore, a violated source AS in the path BC will count as a violated intermediate host AS for all multi-hop paths AD that overlap this particular path. Note that the results presented above correspond to the total number of violating paths and not number of unique exit violating hops.

Fig. 3 presents the number of exit policy violations experienced by each AS traversed by the multi-hop overlay paths. We observe that the number of exit violations is non-uniformly distributed, such that a small fraction of the ASes is violated by a large number of overlay paths.

Lastly, we analyze the relation between the transit violation and exit violation. When we inspect all exit violations that happen in the overlay link AB because of a transit violation at node B, we notice a stronger correlation between the AS experiencing the most exit violations and the node B, rather than with node A. However, a particular transit violating intermediate relay does not have a unique exit violated AS preceding it i.e. there is no one-on-one correspondence between the transit violated AS and the exit violated AS. The number of exit violations per AS, associated with a particular intermediate relay, is often distributed in the same non-uniform manner as in Fig. 3, i.e. Some ASes experience significantly more exit violations compared to others. This observation helps us improve the mitigation strategy in Section IV.

Although this paper only considers the metric of latency at the overlay layer, we expect similar violating behavior with other metrics (e.g., bandwidth) as well, as long as the overlay routing offers substantial improvement over native routing. This can be reasoned by the fact that policy violations are primarily caused by multi-hop overlay paths, which tend to deviate from the direct native route. Thus, the higher the number of multi-hop overlay paths, the higher the extent of policy violations.

It is possible that a certain deviation from the standard end-to-end path may not be objectionable to the AS involved. In our dataset, we observed that 71.3% of the deviating overlay paths (52.6% being of type $E1$) have an AS path that is longer than that of the direct native route. In such cases, we can assert that the exit violations observed are indeed objectionable. However, we are unable to establish with confidence about the seriousness of the other violations. In the rest of the paper, we assume the worst case scenario, where all deviations are serious and are of equal importance to the ASes concerned.

IV. Pricing Scheme for Mitigation

In Section III, we characterized the type and extent of exit policy violations. It is conceivable that the economics and load implications of these violations are objectionable, thereby making it worthwhile to commence filtering overlay traffic in general. We showed in our previous work that the user experience suffers drastically when a large number of ASes enforce these policies by means of filtering[8]. This motivates us to devise a solution that is mutually agreeable to both the native ASes and the overlay service providers (OSP), in the context of exit policy violations.

We devise a pricing-based strategy that monetizes the objections of the violated ASes i.e., we determine the amount an OSP needs to pay the native layer, so as to use an objectionable inter-domain link for its traffic. It could be based on the monetary loss incurred by the native service provider for sending traffic in a different direction. Say an OSP pays C_1 to its host AS X for sending traffic destined to AS Y and causes the host AS to incur an higher expense of C_2 when it performs relaying, then the OSP can pay the overhead of $C_2 - C_1$ to the host AS. This is the only option available in most cases, as enforcing all exit policies will cause the overlay traffic to take the direct native route. Thus, we mitigate the exit violation by:

- Obtaining exit permit from certain ASes to condone any deviation in exit point. This generally incurs a permit fee of E_i for AS 'i' over the lifetime of the overlay[4].

As legitimizing exit violations does not necessarily remove the other transit violations prevalent in the system, we need an integrated solution to resolve all possible violations via this pricing strategy. The general idea behind the cost-sharing approach proposed in [1] is to compensate the native network monetarily, in return for retaining the higher routing gain achieved by the overlay network. In that regard, we undertook the following steps:

- Adding overlay nodes at non-stub ASes, so as to create good overlay paths without transit violations. This incurs a new node fee N_i for AS 'i' over the lifetime of the overlay.
- Obtaining transit permits from certain ASes, so as to allow violating shortest overlay paths to traverse its network. This incurs a permit fee of T_i for AS 'i' over the lifetime of the overlay.

Fig. 4 illustrates the cost-sharing approach adopted for legitimizing the violating overlay path ABC. The OSP has a choice of either adding a new node in AS_{22} and adopting the new path through it, or obtaining a transit permit from AS_{32} and adopting the shortest path. However, neither option eliminates the exit policy violation experienced by AS_{31}. The OSP has fewer alternatives for legitimizing the exit policy violations. This is because any deviation from the direct native route, with an objective of attaining a higher routing gain, will cause more exit violations. Hence, the OSP is required

[4]This fee may be zero if the overlay nodes are already permitted by the end-user agreement to use any desired exit.

to compensate AS_{31} for the exit policy violation in order to retain the routing advantage.

The objective of the pricing strategy is to *determine the exact share of the budget that needs to be spent towards obtaining exit permits, obtaining transit permits and adding new overlay nodes at non-stub networks*. Furthermore, the sequence of these three steps is critical in optimizing the budget usage. This can be attributed to the complex dependency between the two violations, as explained in Section II.

In this section, we describe the exact changes to the cost-sharing approach that will achieve a near-optimal solution. Obtaining the optimal set of changes to make is a hard problem[5]. Hence, we use insights from our analysis in Section III to derive greedy heuristics to obtain a reasonable solution.

A. Greedy Heuristics

In our previous work, we observed that a sequential process of obtaining permits and adding new nodes[1] provides near-optimal tradeoff between budget and path gain. We quantify the overlay routing performance achieved using the *gain* metric, where the gain achieved for a path is defined as:

$$\text{Gain for path AB} = \frac{(\text{Overlay link latency})_{AB} - (\text{Overlay path latency})_{AB}}{(\text{Overlay link latency})_{AB}}$$

The main property of overlay paths that we used to approximate was the *betweenness* value; we define the betweenness of a node as the number of overlay paths transiting through that particular node. Owing to the non-uniform distribution of relay betweenness, we found it worthwhile to obtain transit permits initially and then use the remaining budget for adding new nodes. This ordering was crucial to maximize the improvement in routing performance achieved for a particular budget. The non-italicized portion of Fig. 5 represents the cost-sharing algorithm of [1].

Adding new nodes or permitting a transit typically has the tendency of changing the adopted overlay path, and thereby the exit points. This indicates that the exit violations must be resolved *after each cost-sharing move*. Our analysis of the location of the exit violations showed no evidence that a particular AS is always violated when a certain overlay node is used in that link. However, we did observe a higher correlation between the location of the exit violation and the intermediate relay used. This further corroborates our choice of estimating exit violations after determining the exact intermediate relay to use at each step.

Based on our observations from Fig. 3, we posit that obtaining exit permit from the AS experiencing the most exit violations is the best choice. Nevertheless, if there exists an exit violated AS with very low exit permit fee E, then it would be in our best interest to give a higher preference to obtaining an exit permit from it. We achieve this by normalizing the number of exit violations at each AS by its exit permit fee (unless the cost is zero, for which we just use the absolute value), as done in the approximation algorithm for the set-cover problem[13]. This provides a good tradeoff between budget and path gain.

We present the basic algorithm for resolving all policy violations in Fig. 5. The main factor that determines the amount of budget spent on each operation is the *budget threshold*;

[5]The pricing problem can easily be shown to be NP-hard by performing a reduction to the set-cover problem, which is known to be NP-complete[12].

- Initial total_cost = 0
- For each path i in the set of violating paths
  - For each relay node j in path i, betweenness(j)++
  - Compute Gain(i)
- Sort all relay nodes in the decreasing order of $\frac{\text{betweenness}}{T \text{ for host AS of node}}$
- For each relay node n in that sorted list, while total_cost $<$ budget threshold, do:
  - Obtain transit permit from the host AS i of node n
  - *Compute potential overlay paths after obtaining transit permit*
  - *Determine for each AS k, the* $\frac{\textit{number of exit violations}}{E_k}$
  - *Obtain exit permit from the AS j experiencing the most exit violations*
  - total_cost += $T_i + E_j$
- Sort all transit violating overlay paths in decreasing order of $\frac{\text{path gain}}{N \text{ for upstream provider AS}}$
- For each overlay path i in that sorted list, while total_cost $<$ total budget
  - Add overlay node in AS i that is an upstream provider of the transit violated AS
  - *Compute potential overlay paths after adding new overlay node*
  - *Determine the number of exit violations at each AS*
  - *Determine for each AS k, the* $\frac{\textit{number of exit violations}}{E_k}$
  - *Obtain exit permit from the AS j experiencing the most exit violations*
  - total_cost += $N_i + E_j$

Fig. 5. Greedy scheme for improving path gain without any violations. We have italicized the new modifications that address exit violations.

the exact value adopted depends on the actual topology of the native and overlay network. Nevertheless, we find that a good rule of thumb to decide the value of the threshold is to *determine the budget required to obtain transit permits from ASes that have a betweenness greater than 50% of the maximum betweenness observed in the unrestricted scenario.*

For better understanding of the results, we assume the transit permit fee T_i and the exit permit fee E_i for a particular AS i to be equal to P. Fig. 6 presents the gain observed for each expenditure by the OSP, when the budget threshold was configured at $12 \times P$. From the figure, we observe that the OSP is able to achieve a significant improvement in routing performance that is commensurate with the budget spent. We also see that the increase in gain is sluggish when the budget is low. This is because all ASes in the system filter out violating traffic initially. Obtaining transit permits in that scenario does not cause much change in the gain, until a few critical exit violated ASes are compensated. After crossing that point, the achieved gain increases rapidly with the increase in budget.

Further, we conducted the pricing experiment for a random distribution of the costs N_i, T_i and E_i. The new node fee was uniformly distributed between [$0.5\times N$, $1.5\times N$] and the permit fees between [$0.5\times P$, $1.5\times P$], thus maintaining the average values at N and P. In this scenario, the improvement achieved for a certain budget was similar to the earlier simplified scenario in Fig. 6, showing that the algorithm is more influenced by average costs, rather than the absolute value.

Though there exists no unique AS that needs to be appeased after each cost-sharing step, our greedy sequential approach offers a reasonable improvement in routing performance for the

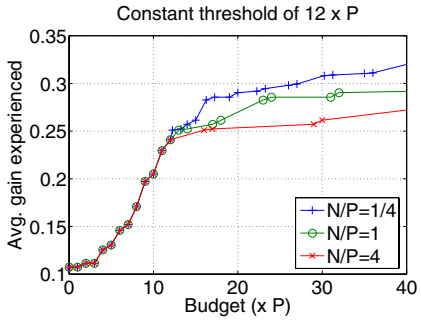

Fig. 6. Pricing-based greedy strategy achieve good gain, when $B_{th} = 12 \times P$. Here, P is equal $\forall$ ASes and N is equal $\forall$ ASes.

OSP. Moreover, we do not notice a case where a cost-sharing step fails to make a difference because the corresponding exit violated AS is not compensated.

Note that it is possible the same AS experiences transit violations in one set of multi-hop overlay paths and exit violations in a different set of multi-hop overlay paths. Nevertheless, we devised it such that the OSP pays individual permit fees for better clarity of our approach.

Although many of the conclusions drawn in this paper may seem limited by the fact that they originate from one realistic overlay testbed, we were able to verify the generality of our approach by inspecting overlay paths in 90 simulated overlay networks. Each of these networks used a large set of real BGP routing tables to determine the AS path at the native layer and a randomly generated inter-domain latency to determine the overlay link latencies[6]. Through these simulations, we asserted to be true, in several scenarios, the key characteristic that exit violations are distributed in a non-uniform manner across multiple ASes. Thus, we are able to apply our pricing strategy in an effective manner over multiple scenarios.

V. Related Work

There are various research efforts that investigate the impact of conflict in objective between the two routing layers, viz. overlay and native. For instance, [14], [15], [16] investigate the interaction between overlay routing and traffic engineering deployed at the native layer. Their general conclusion is that the interaction causes sustained route oscillations and sub-optimal performance for both layers. Our work investigates the conflict between selfish latency-based routing at the overlay layer and inter-domain policies defined at the native layer.

Much of the approach adopted in this paper is similar to our previous work[1]. However, this paper is a direct complement of the previous work, since we address all other policy violations caused by overlay routing in ASes that lie along the legitimate route between two overlay nodes. As mentioned in Section II, the valley-free transit violation occurs at an intermediate host AS, while the exit violations occur at ASes upstream to it. Together, the two classes of violations represent the set of all possible violations in overlay routes.

Our work is similar in spirit to the previous analysis of policy violations at the native layer caused by BGP misconfigurations[17]. Specifically, we investigate the extent to which overlay routes violate the exit preference of each AS; two extremes of this policy are hot-potato and cold-potato routing[18].

[6]Refer to [1] for more information on the simulations.

VI. Concluding Remarks

In this paper, we investigate the concern that overlay routing derives performance advantages by violating native routing policies. Specifically, we investigated the class of exit policy violations caused by overlay paths. We presented a classification of the different types of exit violations possible and their relation to the transit policy violations investigated in an earlier work. By analyzing a case study overlay network built over Planetlab, we characterized the extent, type and location of these exit violations. As the amount of overlay traffic surges, it is conceivable that more networks will start filtering overlay traffic that cause policy violations from an end-to-end perspective. In such a context, it is worthwhile to derive a mutually agreeable solution that achieves good routing performance for the overlay service provider by compensating the native service provider to retain the routing gain. We prescribed a heuristic-based pricing algorithm that obtains a reasonable routing gain for a certain budget. We believe that this approach helps legitimize all native policy violations and allows the benefits obtained by the overlay to be directly related to costs incurred by the overlay service provider.

References

[1] S. Seetharaman and M. Ammar, "Characterizing and Mitigating Inter-domain Policy Violations in Overlay Routes," in *Proceedings of IEEE ICNP*, November 2006.
[2] Y. Chu, S. Rao, and H. Zhang, "A Case for End System Multicast," in *Proceedings of ACM SIGMETRICS*, June 1999.
[3] L. Subramanian, I. Stoica, H. Balakrishnan, and R. Katz, "OverQoS: offering Internet QoS using overlays," in *Proceedings of ACM SIGCOMM*, August 2003.
[4] Lixin Gao, "On inferring autonomous system relationships in the internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, 2001.
[5] "Planetlab," `http://www.planet-lab.org`.
[6] W. B. Norton, "A Business Case for ISP Peering," White Paper, `http://www.equinix.com`, February 2002.
[7] Michael Geist, "Towards a two-tier internet," *BBC News*, vol. `http://news.bbc.co.uk/1/hi/technology/4552138.stm`, December 2005.
[8] P. Grant and J. Drucker, "Phone, Cable Firms Rein In Consumers' Internet Use," Wall Street Journal, Oct, 2005, http://online.wsj.com/article/SB112985651806475197.html.
[9] N. Spring, D. Wetherall, and T. Anderson, "Scriptroute: A facility for distributed Internet measurement," in *4th USENIX Symposium on Internet Technologies and Systems*, March 2003.
[10] Z. Mao, D. Johnson, J. Rexford, J. Wang, and R. Katz, "Scalable and accurate identification of AS-level forwarding paths," in *Proceedings of IEEE INFOCOM*, March 2004.
[11] Beichuan Zhang, Raymond Liu, Daniel Massey, and Lixia Zhang, "Collecting the Internet AS-level topology," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 53–61, 2005.
[12] M. R. Garey and David S. Johnson, *Computer and Intractability: A Guide to the Theory of NP-Completeness.*, W. H. Freeman, 1979.
[13] D. Hochbaum, *Approximation Algorihtms for NP-Hard Problems*, Brooks/Cole Publishing Co., 1996.
[14] L.Qiu, R.Y.Yang, Y.Zhang, and S. Shenker, "On Selfish Routing in Internet-Like Environments," in *Proceedings of ACM SIGCOMM*, August 2003.
[15] Y. Liu, H. Zhang, W. Gong, and D. Towsley, "On the Interaction Between Overlay Routing and Traffic Engineering," in *Proceedings of IEEE INFOCOM*, 2005.
[16] S. Seetharaman and V. Hilt and M. Hofmann and M. Ammar, "Preemptive Strategies to Improve Routing Performance of Native and Overlay Layers," in *To appear in the Proceedings of IEEE INFOCOM*, May 2007.
[17] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *Proceedings of ACM SIGCOMM*, 2002.
[18] Lakshminarayanan Subramanian, Venkata N. Padmanabhan, and Randy H. Katz, "Geographic properties of internet routing," in *Proceedings of the General Track: 2002 USENIX Annual Technical Conference*, 2002, pp. 243–259.