On the Stability of Interdomain Inbound Traffic Engineering

Hao Wang*, Haiyong Xie*, Yang Richard Yang*
*Computer Science Department, Yale University, New Haven, CT 06511

1. INTRODUCTION

In the Internet, ISPs adopt local routing polices to choose routes to achieve objectives such as reducing cost, increasing revenue, reducing latency, and avoiding congestion. Recently, systematic models (*e.g.,* [3]) are proposed to study the stability of path-vector, policy-based interdomain routing. Gao and Rexford [2] prove the surprising result that the constraints on local routing policies due to business considerations can guarantee stability.

Although the preceding stability results are surprisingly pleasant and elegant, practice poses further challenges in analyzing interdomain routing stability. First, the previous studies focus on a specific interdomain route selection algorithm (*e.g.*, the BGP-based greedy route selection algorithm such as SPVP [3]). As a result, factors such as route dampening, which are present in routing practice, are not easily allowed in previous analysis. Although conceptually such factors might not change the conclusions of previous analysis, an analytical framework is still missing. Second, the previous studies focus on local policies which rank only the egress routes; that is, they assume that the local ranking of egress routes at each autonomous system is independent of the inbound traffic pattern of the AS. This independence is justified when the inbound traffic of an AS is relatively constant. However, in practice, the local policies of ASes may involve both the egress routes and the pattern of inbound traffic, introducing unexpected interaction. Specifically, an AS may rank egress routes depending on the pattern of inbound traffic. If this happens, we say that the local policy of the AS is inbound-traffic-dependent, or inbound-dependent for short. One way such inbound-dependent route selection can happen is that the operator of the AS observes traffic demand, and manually reconfigures the local preference values. Such inbound-dependent route selection can also be implemented automatically, with a traffic engineering algorithm based on an estimated traffic demand matrix.

In this paper, we analyze the stability of interdomain routing under the general model that the local preference of an AS depends on not only its egress routes to the destinations but also its inbound traffic pattern. Furthermore, instead of studying a specific route selection algorithm, we study a large class of route selection algorithms which are characterized by their asymptotic behaviors.

Specifically, we first show that the common route selection algorithms can lead to instability due to traffic-route mis-association This instability happens even when all constraints on interdomain routing imposed by business considerations [2] are satisfied, and just a single AS is using such an algorithm. As a remedy, an AS should adopt a route selection algorithm which estimates inbound traffic in such a way that the estimated inbound traffic is truly the result of the chosen egress route.

We then analyze the stability of a network where ASes run any reasonable route selection algorithms which we call *rational route selection algorithms*. The definition of a rational route selection algorithm depends only on the asymptotic behavior of the algorithm. Conducting stability analysis based on the general notion of rational route selection algorithms allows us to prove the stability of a heterogeneous network where different ASes can run different route selection algorithms, so long all of the algorithms are rational. Since the notion of a rational route selection algorithm is defined by its asymptotic behavior, if variations to a route selection algorithm do not change its asymptotic behavior (*e.g.*, non-persistent route dampening), the route selection algorithm is still rational, and thus the stability result still holds.

2. A MOTIVATING EXAMPLE AND THE TRAFFIC MATRIX BASED SCHEME

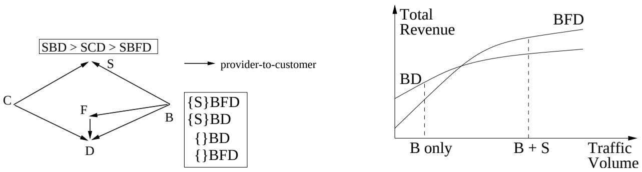

(a) network. (b) revenue function.

Figure 1. An AS with ingress-dependent route ranking table. D is the only destination.

Consider the example shown in Figure 1(a), which is motivated by the increasing usage of multihoming and its potential effects on some transit ISPs. The special feature of this example is that the ranking of AS B, who is one of the two competing providers of S, now depends on *outcomes*, instead of route profiles. An outcome consists of both route selection and ingress traffic pattern for an AS. Specifically, $\{S\}BFD$ denotes the outcome that B uses the route BFD and S sends traffic for destination D through B; $\{\}BD$ denotes the outcome that B uses the route BD and S does not send any traffic through B. This example can well happen in practice. The ranking table of S is constructed according to the standard BGP decision process. As for B, note that B prefers traffic from a customer than no traffic; when S sends traffic through B, the route BFD is preferred than the route BD; otherwise, the route BD is preferred. A potential revenue function that may cause this scenario to happen is shown in Figure 1(b).

2.1. Instability of Traffic Matrix-based Route Selection

A common approach for B to select route is to use a traffic matrix based route selection. Specifically, B could follow the following traffic matrix based route selection protocol: estimate the current ingress traffic matrix, compute the available route such that the chosen route has the highest rank given the demand matrix.

Using this protocol, assume initially S does not use B. Then B first picks BD. Since B chooses BD, S chooses the route SBD, and the traffic from S arrives at B. Since B likes to use BFD when it has high traffic volume, it switches to BFD. Then S chooses SCD, and S no longer uses B. Thus B switches back to BD and we have a loop. This loop shows that the traffic-matrix-based interdomain route selection can be unstable!

The above instability is due to the fact that under the preceding route selection algorithm, B mis-associates the outcomes with its available actions. To choose the optimal route and maintain stability, an AS i needs to correctly associate the outcomes with its actions; that is, the estimated inbound traffic pattern is a result of the chosen egress route.

3. GENERAL RATIONAL ROUTE SELECTION ALGORITHMS

The above instability is due to the fact that under the traffic matrix scheme, B does not keep state to learn the outcome of choosing BD or BFD. When an AS keeps states about the outcomes of choosing different routes and actively seeks to optimize its traffic engineering objectives, we say that the AS is adopting an *active route selection algorithm*.

Among all possible active route selection algorithms, we identify a general class of algorithms which we call *rational route selection algorithms*. Our model is inspired by previous work on adaptive learning [4] and learning on the Internet [1]. The game theoretic models used in the previous work are normal form games. However, BGP is more of an extensive form game than a normal form game, since an intrinsic characteristic of BGP is that route selection of an AS depends on its neighbors.

Intuitively, a reasonable route selection algorithm should not choose route profiles that are shown to be inferior to other available route profiles. Consider the example in Figure 1. At the beginning, B does not know which route (profile) is better, BD or BFD. So it can experiment with each one several times. Later, it may learn that BD always yields a better outcome than BFD does. Thereafter, it is reasonable that B will always choose BD over BFD. The route profile BFD is an example of *overwhelmed* route profiles. Informally, an rational route selection selection algorithm is one where recursively, overwhelmed route profiles are no longer chosen.

Unfortunately, even under the general rational route selection scheme, there are still well-behaved network setups with no stable route selection. Figure 2 shows an example where no rational route selection algorithm can converge to a stable route selection. The setup is constructed to satisfy all standard ISP business relationship constraints so that under previous route selection models [2] there is a unique stable route selection. We observe the following instability when ASes use rational route selection algorithms. When AS A and B choose AD and BFD. The outcome is SAD since S ranks SAD higher than $SBFD$. A has an incentive to change from AD to AED since A ranks $\{S\}AED$ higher than $\{S\}AD$. However, AS B realizes that, it can achieve a better outcome by changing BFD to BD since S will choose SBD over $SAED$. This in turn triggers A to switch from AED back to AD. Thus we end up with A chooses AD and B chooses BFD again, and the process continues forever.

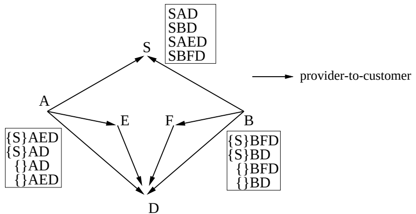

Figure 2. An example with instability. D is the only destination.

Let $r[t]$ denote the network route selection at time t, where $r_i[t]$ is the route selection of AS i. Also, let $\pi_i(r_i, r_{-i})$ be the utility of AS i if the network route selection is r. Furthermore, denote by $A_i(r_{-i})$ the set of routes available to AS i if the route advertisements it receives from its neighbors are given by the corresponding entries in r_{-i}. The instability of the example in Figure 2 under rational route selection scheme is established by the following result:

Theorem 1: Suppose that a sequence of network route selections $\{r(t)\}_{t=0}^{\infty}$ is consistent with rational route selection and that it converges to a stable route selection r^*. Then the following holds for each AS i:

$$\forall r'_i \in A_i(r^*_{-i}), \pi_i(r^*_i, r^*_{-i}) \geq \pi_i(r'_i, r^*_{-i}).$$

An analysis of all of the possible network route selections of the example in Figure 2 shows that no network route selection satisfies the condition in Theorem 1. For this particular example, therefore, no rational route selection algorithm can converge to a stable route selection!

Even other solution concepts from cooperative game do not appear to help for this single example. Such negative results demonstrate the intrinsic challenges of interdomain traffic engineering. Thus global coordinations of preferences of the ISPs are needed to guarantee existence and uniqueness.

Theorem 2: A network running rational route selection algorithms asymptotically converges to the set $U^{\infty}(R)$, where $U(R)$ is the operator that eliminates overwhelmed route selections. Thus, if $U^{\infty}(R)$ is a singleton, the network is guaranteed the existence and uniqueness of stable route selection.

REFERENCES

[1] E. Friedman and S. Shenker. Learning and implementation on the Internet. Working paper. Available at `http://www.orie.cornell.edu/~friedman/pfiles/decent.ps`, 1997.

[2] L. Gao and J. Rexford. Stable Internet routing without global coordination. *IEEE/ACM Transactions on Networking*, 9(6):681–692, Dec. 2001.

[3] T. G. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking*, 10(22):232–243, Apr. 2002.

[4] P. Milgrom and J. Roberts. Adaptive and sophisticated learning in normal form games. *Games and Economic Behaviors*, 3:82–100, 1991.