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 reduce cost, increase revenue, reduce latency, and avoid 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 previous models can capture a wide range of potential interdomain route selection behaviors, they focus only on egress route selection. An implicit assumption of the previous models is that egress route selection and inbound traffic patterns do not interact. However, in practice, egress route selection and inbound traffic patterns could interact. First, the ranking of egress routes by an AS can depend on inbound traffic patterns. A commonly used scheme of egress route selection considering inbound traffic patterns is the traffic-demand-matrix-based scheme. Under this scheme, a traffic demand matrix is first estimated; then the relative ranking of different egress routes is evaluated assuming the given traffic demand matrix. Although this traffic-demand-matrix-based scheme has been shown to be effective, the benefits of the proposed scheme are often evaluated under two assumptions: 1) the route selection of the AS does not change the inbound traffic pattern; and 2) only one AS adopts the proposed scheme. An open question is whether such a scheme can lead to instability when the assumptions are invalid. Second, an AS not only reacts to a given inbound traffic pattern, but may also try to influence it (*i.e.*, a type of inbound interdomain traffic engineering). The large number of prepended prefixes in BGP routes shows that many ASes not only react to given inbound traffic patterns, but also try to influence the patterns. Recently, we have conducted an email survey of ISPs, and the results indicate that the route selection of egress routes of many ISPs depends on inbound traffic patterns; also, ISPs select egress routes to try to influence inbound traffic patterns.

In this paper, we show that the dependency between route selection and inbound traffic patterns can lead to instability under any algorithm of a class of algorithms, even when the constrains of business consideration are followed. This result clearly demonstrates the intrinsic challenges of route selection for interdomain traffic engineering. We also derive a general framework to guarantee convergence and uniqueness of route selection for interdomain traffic engineering.

Hao Wang and Haiyong Xie are graduate students; Y. Richard Yang is a faculty member at Yale University. Email: hao.wang@yale.edu

2. A Motivating Example and the Traffic Matrix Based Scheme

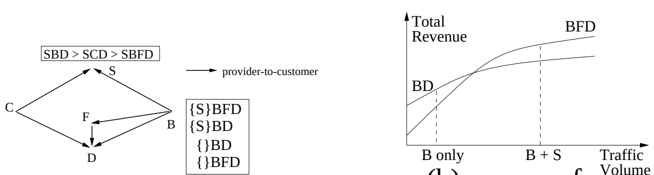

(a) network. (b) revenue func.

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!

3. The Adaptive Route Selection Model

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 *adaptive 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 adaptive route selection algorithm is one where recursively, overwhelmed route profiles are no longer chosen.

Unfortunately, even under the general adaptive route selection scheme, there are still well-behaved network setups with no stable route selection. Figure 2 shows an example where no adaptive 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 adaptive 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.

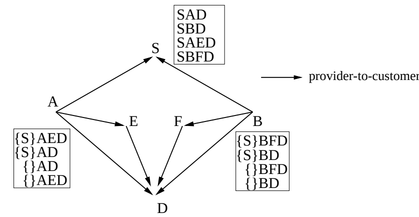

Figure 2. An example with instability. D is the only destination.

The instability of the example in Figure 2 under adaptive 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 adaptive 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 adaptive 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. Specifically, for this game, in each outcome, either A or B can make a change and lead to another outcome. Actually, there is a cycle among the 4 outcomes: $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 1 \rightarrow \cdots$. Therefore, according to the concept of stable set of von Neumann and Morgenstern, the stable set of this example is empty. Thus, there is no solution if we use stable set. According to the concept of bargaining set, the example has a bargaining set which contains more than one element. 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 adaptive route selection algorithms asymptotically converges to the set $U^{\infty}(R)$, where $U(R)$ is the elimination operator. 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.