

On the Stability of Rational, Heterogeneous Interdomain Route Selection*

Hao Wang[†] Haiyong Xie[†] Yang Richard Yang[†]
Li Erran Li[‡] Yanbin Liu[§] Avi Silberschatz[†]

[†] Computer Science Department, Yale University, New Haven, CT 06520

[‡] Networking Research Lab, Bell-labs, Murray Hill, NJ 07974

[§] Department of Computer Sciences, University of Texas, Austin, TX 78712

Abstract

The recent discovery of instability caused by the interaction of local routing policies of multiple ASes has led to extensive research on the subject. However, previous studies analyze stability under a specific route selection algorithm. In this paper, instead of studying a specific route selection algorithm, we study a general class of route selection algorithms which we call rational route selection algorithms. We present a sufficient condition to guarantee routing convergence in a heterogeneous network where each AS runs any rational route selection algorithm. Applying our general results, we study the potential instability of a network where the preference of an AS depends on not only its egress routes to the destinations but also its inbound traffic patterns (i.e., the distribution of incoming traffic from its neighbors). We show that there exist networks which will have persistent route oscillations even when the ASes strictly follow the constraints imposed by business considerations, and adopt any rational route selection algorithms.

1. Introduction

In the Internet, each autonomous system (AS) adopts its own local routing policies to choose interdomain routes to achieve objectives such as cost reduction, revenue maximization, latency reduction, and congestion avoidance. The discovery (e.g., [42]) that the interaction of local routing policies (called local policies for short in this paper) can lead to instability has led to extensive research on the subject lately. By instability in this paper, we mean persistent route oscillations even when the network topology is stable. In particular, researchers [17, 21, 22, 26, 38] study the stability of path-vector, policy-based interdomain routing, and identify conditions to avoid instability. Gao and Rexford [16, 17] prove that the constraints imposed on lo-

cal policies by business considerations can lead to 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 the stability of a homogeneous network where each AS runs the same specific interdomain route selection algorithm (i.e., the BGP-based greedy route selection algorithm such as SPVP [22], where an AS always chooses the best currently available routes). However, with increasing usage of BGP route selection for interdomain traffic engineering, route selection algorithms with more sophisticated strategies are likely to be designed and deployed in the Internet. For instance, Dakdouk *et al.* [4] show an example network where one of the ASes has a route selection strategy which performs strictly better than the greedy strategy. Therefore, given the potential advantage of adopting route selection algorithms that do not use the greedy strategy, different ASes are likely to adopt different route selection algorithms that are suitable for their own objectives. Thus it is necessary to analyze the stability of a *heterogeneous* network where ASes may adopt route selection algorithms that do not use the greedy strategy.

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 AS 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. If this happens, we say that the local policy of the AS depends on the inbound traffic pattern, or inbound traffic for short. We also say that the local policy of the AS is inbound-traffic-dependent, or inbound-dependent for short. Later in Section 5, we will show an example network where one of the ASes ranks egress routes depending on the pattern of inbound traffic. Such inbound-dependent local policies can be implemented automatically with a traffic engineering algorithm based on an estimated traffic demand matrix. In the last few years, several traffic-demand-matrix-based traffic engineering algorithms have been proposed (e.g., [3, 19]). Although such route selection algorithms

*Hao Wang is supported by NSF grants ANI-0207399 and CNS-0435201. Haiyong Xie is supported by NSF grants ANI-0238038 and CNS-0435201. Yang Richard Yang is supported in part by NSF grants ANI-0238038 and CNS-0435201.

have been shown to be effective, the evaluations often assume that the inbound traffic is constant (*e.g.*, the route selection of the AS does not change the inbound traffic), whereas the inbound traffic is likely to change with the chosen egress routes, introducing unexpected interaction. Thus it is necessary to analyze the stability of route selection algorithms implementing local policies that take into account inbound traffic patterns.

In this paper, we analyze the stability of interdomain routing in a heterogeneous network where ASes run any one of a class of route selection algorithms. Informally, the class of algorithms we study are those that, asymptotically, for the given network, will not choose routes that are known to be inferior to other available routes. Since we are modeling the route selection behaviors of self-optimizing ASes, it will be “unjustified” or “irrational” for a self-optimizing AS to eventually choose an inferior route when there are other available, better routes; thus we call the class of algorithms we study *rational route selection algorithms*.

There are several advantages in conducting stability analysis based on the general notion of rational route selection algorithms. First, it allows us to establish more general positive results: 1) it allows us to prove the stability of a heterogeneous network where different ASes can run different route selection algorithms, as long as the algorithm chosen by each AS is rational for the given network; 2) 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 experimentation), the route selection algorithm is still rational, and thus the stability result still holds. Second, it allows us to establish more general negative results; for example, if we show that a network is unstable under *any* rational route selection algorithms, it is more general than to show that a network is unstable under a specific route selection algorithm.

In particular, we derive a sufficient condition to guarantee routing convergence under the general model that the ASes are running heterogeneous rational route selection algorithms. This condition applies to any network where the route selection algorithms of the ASes are rational route selection algorithms.

Applying our general results, we study the potential instability of inbound-dependent route selection. We first show that the common route selection algorithms of simply choosing the best routes according to the traffic demand matrix of the *preceding period* could lead to instability, when the route selection of an AS can change its inbound traffic pattern. This instability happens even when all constraints on interdomain routing imposed by business considerations [17] are satisfied, and just a single AS is using such an algorithm. We say that such instability is caused by traffic-route mis-association, and it is an example of instability caused by route selection algorithms.

Although there is a simple rational route selection algorithm to handle the preceding scenario, we also show that

there exist networks which can have persistent route oscillations even when the local policy of each AS follows the constraints imposed by business considerations, and can adopt *any* one of the rational route selection algorithms. This result clearly demonstrates the intrinsic challenges of inbound-dependent route selection for interdomain routing.

The rest of this paper is organized as follows. In Section 2, we discuss related work. In Section 3, we define the class of rational route selection algorithms. In Section 4, we present a sufficient condition to guarantee convergence of a network running rational route selection algorithms. In Section 5, we show that the traffic-demand-matrix-based route selection algorithms can lead to routing instability. In Section 6, we show an example network which is unstable under any rational route selection algorithms. Our conclusion and future work are in Section 7.

2. Related Work

There is a large body of literature on interdomain route selection. Researchers have conducted extensive evaluations (*e.g.*, [5, 10, 20, 29, 30, 42]) and theoretical analysis (*e.g.*, [17, 21, 22, 25, 26, 38]) on the stability of BGP route selection. In particular, Griffin *et al.* [22] show that “policy disputes” can cause persistent route oscillations. Griffin and Wilfong [23] then propose a protocol called SPVP3 that can detect oscillations caused by policy dispute at run time using “path history”. Gao and Rexford [16, 17] observe that, if every AS considers each of its neighbors as either a customer, a provider, or a peer, and obeys certain local constraints on preference and export policies, then BGP is guaranteed to converge. Generalizing the above commercial relationships of ISPs to a class-based system, Jaggar and Ramachandran [25] show that a global constraint that guarantees convergence can be enforced by a distributed algorithm. Mao *et al.* [34] also describe a mechanism to damp route oscillations. A major difference between our study and the preceding studies is that we analyze the stability of interdomain routing in a heterogeneous network where ASes run any rational route selection algorithms, instead of a specific algorithm (*i.e.*, the greedy route selection algorithm). Also, we study the dependency of route selection on inbound traffic, an important factor which has not been addressed before.

In order to investigate the existence and nonexistence of stable route selection for a heterogeneous network running rational route selection algorithms, we adopt a general, rational, learning model. This model is motivated by general game-theoretical, rational algorithms (*e.g.*, adaptive and sophisticated learning algorithms [35]). In particular, our model is inspired by the adaptive learning model of Milgrom and Roberts [35], and the reasonable learning model of Friedman and Shenker [13–15].

The interaction of interdomain routing and inbound traffic starts to receive some attention lately [18, 43]. How-

ever, the focus of previous studies is on prepending. In [43], Wang *et al.* characterize the stability of inbound-dependent route selection. However, their study focuses on prepending and their specific algorithm. Unlike [43], we focus on route selection, since we feel that the effects of prepending cannot be guaranteed since an AS can choose to ignore the effects of prepending.

The study on traffic engineering has traditionally been focused on intra-domain (for a good survey, please see [11, 12]). There is an increasing interest in tuning BGP attributes for interdomain traffic engineering [5, 37]. However, most of the previous work focuses on egress route selection for either a single AS (*e.g.*, [3, 6, 19]), or between two neighboring ASes. In particular, researchers have conducted extensive theoretical analysis (*e.g.*, [27]) and experimental evaluations (*e.g.*, [40, 41]) of hot-potato routing, which is a scheme of exit route selection between two ASes. Recently, Wang *et al.* [24] study general interdomain egress traffic engineering and identify sufficient conditions for convergence; however, it still focuses only on the greedy route selection algorithm, and considers egress routes only.

Another line of related research is the extensions/alternatives to BGP (*e.g.*, the mechanism-design approach by Feigenbaum *et al.* [7–9], the negotiation protocol by Mahajan *et al.* [31–33], the BGP pricing approach by Afergan and Wroclawski [1], and the Hybrid Link-state Path-vector (HLP) approach of Subramanian *et al.* [39]). The objective of our study is to investigate the intrinsic instability of interdomain routing so that the extensions can guarantee stability under all scenarios.

3. General Rational Route Selection Algorithms

Previous studies on the stability of interdomain routing focus on one specific interdomain route selection algorithm — the greedy route selection algorithm. However, BGP route selection has increasingly been used by ASes to achieve a diverse set of interdomain traffic engineering objectives. For ASes with certain objectives, the greedy route selection algorithm is no longer the best choice. Figure 1 (in the same spirit as the one in Dakdouk *et al.* [4]) shows such an example, where A has a route selection strategy that performs strictly better than the greedy strategy. In more detail, A and B 's traffic engineering objectives require joint ranking of routes to two destinations D_1 and D_2 . The ranking tables are shown in the two boxes. Suppose both A and B start with empty routes, B uses the greedy strategy, and A makes announcement first. Under the greedy strategy, A will select and announce $(AD_1, AG_1G_2D_2)$. This will lead B to select and announce (BAD_1, BFD_2) , and the network becomes stable. However, if A selects and announces inferior routes $(AE_1D_1, AG_1G_2D_2)$ to B , B will select and announce $(BH_1H_2D_1, BD_2)$ to A . This enables A to select the highest ranked routes (AE_2D_1, ABD_2) as

its stable route selection, making this route selection strategy better for A than the greedy strategy. Thus, it is important to analyze a heterogeneous network where different ASes may run different route selection algorithms, not a homogeneous network where all ASes run a single, specific algorithm, *e.g.*, the greedy route selection algorithm.

Specifically, in this paper, we study a class of route selection algorithms we call rational route selection algorithms. Intuitively, a rational route selection algorithm is one which, asymptotically, will not choose routes that are known to be inferior to some other available routes. The concept of rational route selection algorithms is motivated by previous work on adaptive learning [35] and learning on the Internet [14]. The models used in the previous game theoretical studies are normal form games. However, interdomain route selection is more of an extensive form game than a normal form game, since an intrinsic characteristic of interdomain route selection is that the available routes of an AS depend on those exported by its neighbors. In this paper, we shall explicitly model this dependency. In the sequel, we shall present the network model and formalize our intuitive notion of rational route selection algorithms and explore the implications.

3.1. Network Model

The network topology is represented by a simple, undirected graph $G = (V, E)$, where $V = \{1, \dots, N\}$ is the set of ASes and E is the set of interdomain links.

A path in G is either the empty path, denoted by ϵ , or a sequence of ASes $(v_k, v_{k-1}, \dots, v_1, v_0)$, where $k \geq 0$ is the length of the path, such that $(v_i, v_{i-1}) \in E$ for $i = k, k-1, \dots, 1$. Note that if $k = 0$, then (v_0) represents the trivial path from v_0 to itself. Each nonempty path $P = (v_k, v_{k-1}, \dots, v_1, v_0)$ has a direction from v_k to v_0 . If P and Q are two nonempty paths such that the first AS in Q is the same as the last AS in P , then PQ denotes the path formed by the *concatenation* of these two paths. We extend this with the convention that $\epsilon P = P\epsilon = P$ for any path P .

We denote by R the set of all paths in G . For each $i \in V$, we denote by $R_{i \rightarrow}$ the set of paths originating from i , and by $R_{\rightarrow i}$ the set of paths terminating at i . Also, for any $i, j \in V$, $R_{i \rightarrow j} = R_{i \rightarrow} \cap R_{\rightarrow j}$ denotes the set of paths from i to j .

Suppose i and j are two neighboring ASes. As a path P is exported from j and imported into i , it undergoes two transformations. First, $P_1 = \text{export}(i, j, P)$ represents the application of export policies of j to P , which includes possibly prepending j multiple times to P or filtering out P altogether ($P_1 = \epsilon$). Second, $P_2 = \text{import}(i, j, P_1)$ represents the application of import policies of i to P_1 . In particular, import policies at i will filter out any path that contains i itself. The collective effects of these transformations can be represented by the *peering transformation*, $\text{pt}(i, j, P)$,




Figure 1. Illustration of a non-greedy route selection strategy.

defined as

$$\text{pt}(i, j, P) = \begin{cases} \text{import}(i, j, \text{export}(i, j, P)) & \text{if } (i, j) \in E, \\ \epsilon & \text{otherwise.} \end{cases}$$

The peering transformation represents the import/export policies of all ASes in the network. Note that in the above definition, we extend the domain of pt to all pairs of ASes by setting $\text{pt}(i, j, P) = \epsilon$ if i and j are not neighbors.

Each AS $i \in V$ has a set $\mathcal{D}_i \subseteq V$ of destinations, and attempts to establish a path to each destination $j \in \mathcal{D}_i$. A *network route selection* is a function r that maps each pair of ASes $i \in V$ and $j \in \mathcal{D}_i$ to a path $r(i, j) \in R_{i \rightarrow j}$. We interpret $r(i, j) = \epsilon$ to mean that i is not assigned a path to j . We denote by \mathcal{R} the set of all possible network route selections. When we restrict our attention to the route selection of AS i alone, we shall refer to the restriction of r on i and \mathcal{D}_i as the *route profile* for AS i , denoted by r_i . We denote by \mathcal{R}_i the set of all possible route profiles for AS i . Note that in the above definition, we do *not* require the routes in a network route selection to be consistent; that is, if $r_i(k) = (i, j)P$, it is not necessary that $r_j(k) = P$.

The above definitions lead to useful equivalent representations of network route selections and route profiles. First, a network route selection r can be represented as $r = (r_i, r_{-i})$, where $r_{-i} = (r_j)_{j \neq i}$ denotes the *combined route profiles* of all ASes except i . The route profile of AS $j \neq i$ in r_{-i} is denoted by $(r_{-i})_j$. We denote by \mathcal{R}_{-i} the set of all possible combined route profiles of all ASes except i ; that is, $\mathcal{R}_{-i} = \{r_{-i} | (r_{-i})_j \in \mathcal{R}_j, \forall j \neq i\}$. Second, network route selections and (combined) route profiles can be treated as sets of paths. Specifically, a network route selection r , a route profile r_i and a combined route profile r_{-i} are equivalent to the sets of paths $\{r(i, j) | i \in V, j \in \mathcal{D}_i\}$, $\{r_i(j) | j \in \mathcal{D}_i\}$, and $\{(r_{-i})_j(k) | k \in \mathcal{D}_j, j \neq i\}$, respectively. This equivalent representation is particularly convenient in some operators defined on sets of paths. For example, we can simply use r_{-i} as an argument to such an operator, where actually the argument is $\{(r_{-i})_j(k) | k \in \mathcal{D}_j, j \neq i\}$.

An intrinsic characteristic of path vector protocols such as BGP is that there are dependencies among route selections of ASes. Specifically, the route profiles available to

i depend on the route advertisements it receives from its neighbors, which in turn depend on route selections of these neighbors. To capture this dependency, we define two operators C_i and A_i for each AS i as follows. For a set of paths $\mathcal{P} \subseteq R$, let

$$C_i(\mathcal{P}) = \{(i, j) \text{pt}(i, j, P) | P \in \mathcal{P} \cap R_{j \rightarrow}\} \quad (1)$$

$$A_i(\mathcal{P}) = \{r_i \in \mathcal{R}_i | r_i(k) \in C_i(\mathcal{P}) \cup \{\epsilon\}, \forall k \in \mathcal{D}_i\} \quad (2)$$

Intuitively, if \mathcal{P} is the set of routes exported by i 's neighbors, then $C_i(\mathcal{P})$ is the set of routes available to i in its routing cache, and $A_i(\mathcal{P})$ is the set of route profiles that i can possibly choose from this routing cache. Note that AS i can always choose the empty path to any $k \in \mathcal{D}_i$ regardless of $C_i(\mathcal{P})$.

The route selection objective of AS i (*i.e.*, its local preference) is represented by a utility function $u_i(r_i, r_{-i})$, which evaluates the payoff of the current network route selection r for i . Note that since we allow the utility of i to depend on not only i 's route, but also all other ASes' routes, it captures inbound-dependent route selection.

3.2. Algorithm Model

As is mentioned at the beginning of this section, we want to analyze the stability of a heterogeneous network for a wide range of potential route selection algorithms. The only condition we impose on a route selection algorithm is that, asymptotically, it will not choose routes that are known to be inferior to some other available routes. Since we are modeling self-optimizing ASes, we feel that this is a very generic characterization of such ASes. To capture such generic behaviors, we avoid any detailed specification of how the ASes actually select route profiles. Instead, we focus on the sequence of network route selections over time, and define the class of algorithms we consider by identifying the general properties of the sequences generated by the route selection algorithms.

We assume that there is a set of times $T = \{0, 1, 2, \dots\}$ at which one or more ASes in the network change their route profiles. The elements of T should be viewed as the indices of the sequence of physical times at which these changes

take place. At time t , the selected route profile of AS i is $r_i[t]$, and the network route selection is $r[t] = (r_i[t])_{i \in V}$. The sequence of network route selections is, therefore, $\{r[t]\}_{t=0}^{\infty}$.

Given a set $H \subseteq \mathcal{R}$ of network route selections, we define the *projection* of H onto \mathcal{R}_i as

$$H_i = \{r_i \in \mathcal{R}_i | r \in H\}. \quad (3)$$

Accordingly, we define the *product* set H_{-i} as

$$H_{-i} = \{r_{-i} \in \mathcal{R}_{-i} | (r_{-i})_j \in H_j, \forall j \neq i\}. \quad (4)$$

The set H_{-i} represents all possible combined route profiles of all ASes except i , where AS j 's route profile is drawn from H_j for all $j \neq i$. Also, let

$$A_i(H_{-i}) = \bigcup_{r_{-i} \in H_{-i}} A_i(r_{-i}). \quad (5)$$

Recall that in the above definition, $A_i(r_{-i})$ actually means $A_i(\{(r_{-i})_j(k) | k \in \mathcal{D}_j, j \neq i\})$.

In order to get some intuition about the definition of rational route selection algorithms, suppose that AS i has observed a history H of network route selections. If this history is long enough for AS i to believe that it has observed all possible route profiles that will be used by each other AS in the future, AS i will expect that each other AS j will select route profiles in H_j . It is reasonable, therefore, for i to believe that the combined route profiles of the other ASes will belong to the set H_{-i} , hence that the route profiles possibly available to it will belong to the set $A_i(H_{-i})$. However, not all possibly available route profiles in $A_i(H_{-i})$ are worth considering. If there exist two route profiles $r_i, r'_i \in A_i(H_{-i})$, such that the following two conditions hold:

- C1. whenever r_i is available, r'_i is also available;
- C2. choosing r'_i always yields strictly higher payoff than r_i ,

then it would be “unjustified” or “irrational” for i to choose r_i . This is because by C1, r'_i can always be chosen instead of r_i ; and by C2, choosing r'_i always yields strictly higher payoff than choosing r_i . In this case, r_i is said to be *overwhelmed* by r'_i with respect to H , and is called an overwhelmed route profile. A route profile $r_i \in A_i(H_{-i})$ that is not overwhelmed by any other $r'_i \in A_i(H_{-i})$ with respect to H is called an *unoverwhelmed* route profile. If we use $U_i(H)$ to denote the set of unoverwhelmed route profiles of AS i with respect to H , then the definition of $U_i(H)$ requires taking negations of the above two conditions. Formally, we define the following operator $U : 2^{\mathcal{R}} \mapsto 2^{\mathcal{R}}$:

Definition 1 Given $H \subseteq \mathcal{R}$, let

$$\begin{aligned} U_i(H) &= \{r_i \in A_i(H_{-i}) | \forall r'_i \in A_i(H_{-i}), P1 \vee P2, \\ &\text{where} \\ &\quad (P1) \exists r_{-i} \in H_{-i}, \text{ such that} \\ &\quad \quad r_i \in A_i(r_{-i}), r'_i \notin A_i(r_{-i}), \\ &\quad (P2) \exists r_{-i}, r'_{-i} \in H_{-i}, \text{ such that} \\ &\quad \quad r_i \in A_i(r_{-i}), r'_i \in A_i(r'_{-i}), \\ &\quad \quad u_i(r_i, r_{-i}) \geq u_i(r'_i, r'_{-i}), \\ U(H) &= \{r \in \mathcal{R} | r_i \in U_i(H)\}. \end{aligned}$$

The two predicates P1 and P2 in Definition 1 are negations of conditions C1 and C2, respectively. If AS i believes that other ASes will select route profiles in H_{-i} , then it would be “irrational” for AS i to choose any route profile not in $U_i(H)$, since every such route profile is guaranteed to be overwhelmed by some other route profile in $U_i(H)$. $U_i(H)$ thus formalizes our notion of the set of unoverwhelmed route profiles for AS i when each other AS j is limited to route profiles in H_j .

With the definition of unoverwhelmed route profiles, we can now formalize our intuitive notion of “rational route selection” as such that, asymptotically, will not choose overwhelmed route profiles. Formally,

Definition 2 $\{r_i[t] | t \in T\}$ is consistent with rational route selection if, for all t' , there exists $t'' > t'$ such that for all $t > t''$, $r_i[t] \in U_i(\{r[s] | t' \leq s < t\})$. $\{r[t] | t \in T\}$ is consistent with rational route selection if each $\{r_i[t] | t \in T\}$ has this property.

Remark 1 The sequence $\{r[t] | t \in T\}$ is determined by many factors, and thus whether a route selection algorithm used by an AS is rational or not also depends on these factors, which include, but are not limited to, network topology, local policies of ASes in the network, and route selection algorithms used by other ASes. This allows more algorithms to be classified as rational. For example, we will show later that the BGP-based greedy route selection algorithm is rational in a particular type of networks. Also note that there are no requirements on the route selection behaviors of the ASes for the finite period of time from t' to t'' in Definition 2. This allows for an AS to use non-greedy strategies such as the one discussed for the example in Figure 1.

3.3. Rational Route Selection Algorithms

The preceding definition of rational route selection is generic and does not specify how ASes actually select route profiles. Thus, it allows both centralized and distributed implementations. An example centralized implementation can be as follows. Each AS sends its utility function (policies) to a trusted third party. The third party then applies the operator U to compute for each AS a routing schedule (namely what route each AS should adopt at what time).¹

¹This approach can be made possible by the availability of a public database publishing AS routing policies. The ASes should be semi-honest in that they do not manipulate their policies when reporting their policies.

The above implementation requires complete information, due to its generality. As we limit the generality, there can be efficient implementations without requiring complete information, in a distributed setting. In particular, we will analyze the *standard BGP route selection protocol* as it is used in interdomain route selection, and show that it is a distributed rational route selection algorithm. By the standard BGP route selection protocol, we mean essentially the simple path vector protocol (SPVP) as defined in Fig. 5 of [22], extended to the case of joint multiple-destination route selection, when some mild conditions are satisfied. We will show that the asymptotic best-response nature of BGP makes it a rational route selection algorithm, when the ranking of egress routes of an AS depends on the its own egress routes only.

Specifically, we have the following result:

Theorem 1 *The BGP protocol is consistent with rational route selection, if the following conditions are satisfied:*

- A1. *BGP update messages between neighboring ASes are delivered reliably in FIFO order, and have bounded delay;*
- A2. *Each AS sends out BGP update messages in bounded time after it updates its route profile;*
- A3. *Each BGP update message is processed immediately.*

Proof: Let the sequence of network route selections be $\{r[t]\}_{t=0}^{\infty}$.

Consider an arbitrary AS i . Let \mathcal{N}_i be the set of neighbors of i . For any $j \in \mathcal{N}_i$, let $r_j[\tau_j^i(t)]$ be the latest route profile of j such that an update message has been sent to i with this route profile. Thus $C_i(r_j[\tau_j^i(t)])$ is the set of paths in i 's routing cache learned from j at time t . The set of route profiles available to i is therefore $A_i(\{r_j[\tau_j^i(t)] | j \in \mathcal{N}_i\})$. Assumptions A1 and A2 imply that there exists t_d such that at any time t , for any neighbor j of i , $\tau_j^i(t) \geq t - t_d$.

Although i may not know $r_{-i}[t]$, the payoff $u_i(r_i, r_{-i})$ is only a function of r_i . (Recall that we consider only egress route selection in this case.) The BGP protocol, together with Assumption A3, implies that at any time t

$$r_i[t] = \arg \max_{r_i \in A_i(\{r_j[\tau_j^i(t)] | j \in \mathcal{N}_i\})} u_i(r_i, r_{-i}[t]). \quad (6)$$

We shall prove the theorem by showing that $t'' = t' + t_d$ satisfies Definition 2. In fact, for any $t > t''$, let $H = \{r[s] | t' \leq s < t\}$. For any neighbor j of i , we have $\tau_j^i(t) \geq t - t_d \geq t'$, thus $r_j[\tau_j^i(t)] \in H_j$. Therefore, there exists $r_{-i} \in H_{-i}$ such that $r_j[\tau_j^i(t)] = (r_{-i})_j$. We shall show that $r_i[t] \in U_i(H)$. We have that $r_i[t] \in A_i(r_{-i}) \subseteq A_i(H_{-i})$. For any $r'_i \in A_i(H_{-i})$, if predicate P1 does not hold, then $r'_i \in A_i(r_{-i})$, which, together with Equation (6), implies that $u_i(r_i[t], r_{-i}[t]) \geq u_i(r'_i, r_{-i}[t])$. It follows that $r_i[t] \in U_i(H)$. ■

Remark 2 *These three assumptions of the theorem should be valid under normal network operations.*

Remark 3 *Note that in Definition 2, AS i is not required to know the route selections $r_{-i}[t]$ of the other ASes. AS i may not even know the sequence of times T and its set of all possible route profiles \mathcal{R}_i . In addition, the definition says nothing about the routing cache of i . The $r_{-i} \in H_{-i}$ used in Definition 1 may have never appeared in i 's routing cache from time t' up to t . Moreover, at some time t , $r[t]$ may not even be consistent. All that is required is that the exhibited sequences of route selections $r_i[t]$ and $r[t]$ satisfy the requirement in the definition. The preceding theorem is an example clarifying this subtlety.*

4. A Sufficient Condition to Guarantee Convergence of Rational Route Selection Algorithms

Given the definition of rational route selection algorithms, in this section, we derive a sufficient condition to guarantee stability. The advantage of deriving a sufficient condition using the general notion of rational route selection algorithms is that we then only need to consider the asymptotic behaviors of route selection algorithms, allowing variations such as limited route experimentation.

We first define the notion of stable route selection.

Definition 3 *A network consisting of ASes each of which is running a rational route selection algorithm has a stable route selection, if the route selection of each AS has a single route profile, as time goes to infinite. Formally, the network has a stable route selection if $\{r[t]\}_{t=0}^{\infty}$ converges.*

Remark 4 *In the above definition, we require that, in a stable route selection, the route selection of each AS be a "pure" routing decision. We do not allow "mixed" strategies [36], since mixed strategies involve frequent route fluctuations, and are thus not desirable as "stable" solutions for global interdomain routing.*

We first observe the following important property of the operator U :

Lemma 2 *The operator U is monotone: If $P, Q \subseteq \mathcal{R}$ and $P \subseteq Q$, then $U(P) \subseteq U(Q)$.*

Proof: It suffices to show that $U_i(P) \subseteq U_i(Q)$ for an arbitrary i .

Suppose $r_i \in U_i(P)$. We first notice that, since the operator A_i as defined in (2) is monotone, $r_i \in A_i(P_{-i})$ implies $r_i \in A_i(Q_{-i})$. To prove $r_i \in U_i(Q)$, we only need to show that, for any $r'_i \in A_i(Q_{-i})$, at least one of the two predicates P1 and P2, which are defined in Definition 1, holds. We distinguish the following two cases:

1. $r'_i \in A_i(P_{-i})$. In this case, the fact that $r_i \in U_i(P)$ implies that at least one of the two predicates P1 and P2 holds.
2. $r'_i \notin A_i(P_{-i})$. This case happens only if $\forall r_{-i} \in P_{-i}, r'_i \notin A_i(r_{-i})$. Thus predicate P1 holds in this case.

■

Let $U^{(k)}(\mathcal{R})$ denote the k -th iteration of the operator U on \mathcal{R} , for $k = 0, 1, \dots$, with $U^{(0)}(\mathcal{R}) = \mathcal{R}$. We now observe that sequences consistent with rational route selection share some common asymptotic properties:

Theorem 3 *If $\{r[t]|t \in T\}$ is consistent with rational route selection, then for each k , there exists $t_k \in T$ such that, for all $t \in T$ with $t \geq t_k$, $r[t] \in U^{(k)}(\mathcal{R})$.*

Proof: For $k = 0$, the conclusion holds trivially (choosing $t_0 = 0$) since for all t , $r[t] \in \mathcal{R} = U^{(0)}(\mathcal{R})$.

Suppose the conclusion holds for $k - 1$. Then, there is a t_{k-1} such that for all $t \geq t_{k-1}$, $\{r[s]|t_{k-1} \leq s \leq t\} \subseteq U^{(k-1)}(\mathcal{R})$. Since $\{r[t]|t \in T\}$ is consistent with rational route selection, in Definition 2 we may choose $t' = t_{k-1}$ and we may take $t_k > \max(t'', t_{k-1})$. Therefore, for all $t \geq t_k$, we have that $r[t] \in U(\{r[s]|t_{k-1} \leq s < t\})$. By the induction hypothesis and Lemma 2, $U(\{r[s]|t_{k-1} \leq s < t\}) \subseteq U(U^{(k-1)}(\mathcal{R})) = U^{(k)}(\mathcal{R})$. Thus, for all $t \geq t_k$, $r[t] \in U^{(k)}(\mathcal{R})$. ■

By Theorem 3, when the serially unoverwhelmed set $U^\infty(\mathcal{R}) = \bigcap_{k=1}^\infty U^{(k)}(\mathcal{R})$ is small, one can predict with precision the asymptotic behavior of a sequence of network route selections. In particular, if $U^\infty(\mathcal{R})$ is a singleton, Theorem 3 immediately implies that the sequence will always converge to a unique network route selection. We therefore extend similar results in the context of strategic learning game [35] and learning in the Internet [14] to our route selection context.

Proposition 4 *The network route selection of a network consisting of ASes running rational route selection algorithms asymptotically lie in the set $U^\infty(\mathcal{R})$. Thus, if $U^\infty(\mathcal{R})$ is a singleton, the network is guaranteed the existence and uniqueness of stable route selection.*

One way to guarantee that $U^\infty(\mathcal{R})$ is a singleton is the existence of a *sequentially dominant route selection*.

Definition 4 *A network has a sequentially dominant route selection (SDRS) if there is a partial order of the ASes, with the destination being the first one, such that given the route selection of the ASes before i in this partial order, the best route selection of i is determined, independent of the route selection of those after i .*

If a network has an SDRS, all routes other than the unique solution are not in the unoverwhelmed set. As such,

$U^\infty(\mathcal{R})$ is a singleton. The convergence of such networks under any rational route selection algorithms, therefore, follows immediately from Theorem 1 and Proposition 4. Note that the existence of SDRS can be checked in polynomial time.

As an application of the preceding results, we derive a sufficient condition to guarantee routing convergence in a heterogeneous network where each AS runs any rational route selection algorithm, and its egress route selection satisfies the constraints imposed by business considerations [17].

Theorem 5 *Assume a network where each AS runs any rational route selection algorithm, and selects egress routes independent of inbound traffic. Assume that 1) there is no provider-customer loop in the network; and 2) each AS adopts the typical export policy and the standard joint-route preference [44]. Then $U^\infty(\mathcal{R})$ is a singleton; that is, the network is guaranteed to converge to the unique stable route.*

Proof: (sketch) When the conditions of the theorem are satisfied, we can use an induction proof to show the existence of an SDRS. Therefore, the network is guaranteed to converge to the unique stable route. ■

Remark 5 *The preceding convergence result is more general than that proved in previous studies in that it is not limited to just homogeneous networks where each AS has to run the greedy, best-response BGP algorithm. Other actions, such as non-persistent experimentation are allowed.*

5. Inbound-dependent Route Selection: Traffic-Demand-Matrix-Based Algorithms

5.1. A Motivating Example

Starting from this section, we apply our general framework of rational route selection algorithms to study the stability of a network when ASes may adopt general local policies that take into account inbound traffic patterns.




Figure 2. The ranking of egress routes at B depends on inbound traffic. S is the source, and D is the destination.

We start with an example shown in Figure 2. The example is motivated by the increasing usage of multihoming and its potential effects on some transit ISPs. A special feature

of this example network is that the ranking of egress routes at B , who is one of the two competing transit providers of source S , depends on its inbound traffic. For generality, we say that B ranks *outcomes*, instead of just egress routes. An outcome consists of both an egress route and ingress traffic pattern. For generality, we assume a ranking table at each AS, which lists, in decreasing order, all of the potential outcomes. Note that in practice, a ranking table can be implemented, compactly, by an objective or utility function. Specifically, $\{S\}BFD$ denotes the outcome that B uses the egress 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 network does not appear to be a pathological case and can well happen in practice. S is a multihomed network with two providers C and B to improve reliability. The ranking table of S is constructed according to the standard BGP decision process: S prefers routes with small AS-hop counts; for two routes with the same AS-hop count, it uses the next-hop ID to break the tie. As for B , when traffic volume is high (*i.e.*, when S uses B as its transit provider), B selects BFD over BD ; on the other hand, when traffic volume is low (*i.e.*, when S does not use B as its transit provider), B chooses BD over BFD . A potential revenue function that may cause this scenario to happen is shown in Figure 3; that is, BFD is more profitable for B when the traffic volume is high, while BD is more profitable for B when the traffic volume is low. Note that it is possible to reverse the provider-customer relationship of the AS pairs, CD , FD , BF , and BD . Then the preference of B can be justified by cost instead of revenue.




Figure 3. A revenue function justifying the route selection behavior of B in Figure 2. “ B only” denotes the traffic volume when S does not use B as its transit provider; and $B + S$ denotes that when S uses B .

5.2. Instability of a Traffic-Demand-Matrix-Based Greedy Route Selection Scheme

A common approach for B to implementing inbound-dependent route selection is to use a traffic-demand-matrix-based algorithm (*e.g.*, [3, 19]). The basic structure of such an algorithm is that time is divided into multiple periods.

During each time period, the algorithm measures the traffic demand matrix. At the end of each time period, the algorithm computes and installs the optimal route selection for the next period.

In particular, B could implement a route selection algorithm using the greedy strategy as follows. During each time period n , B estimates total traffic demand to destination D ; At the end of time period n , B computes the optimal route selection (BFD or BD), based on the measured inbound traffic demand and its traffic engineering objectives. B then installs the optimal route selection at the beginning of time period $n + 1$. As we have discussed in the introduction, this algorithm can be implemented either by a network operator manually, which will operate at a longer time scale, or by a traffic engineering program, which will operate at a much faster speed.

However, this traffic-demand-matrix-based greedy route selection algorithm is not a rational route selection algorithm (which we will show later). It will also cause routing instability in the example network. To see this, assume that B initially chooses egress route BD . B exports BD to S ; therefore, S chooses SBD over SCD , and the traffic from S to D goes through B . However, given this high inbound traffic demand, B prefers BFD over BD ; thus B switches its route selection to BFD and exports to S . This change of egress route causes S to choose SCD over SBD , and thus traffic of S no longer goes through B . Given that now the inbound traffic is low, B switches back to route selection BD , since it prefers BD over BFD at low traffic. Thus, we have obtained persistent route oscillations².

The above instability is due to the fact that under the preceding traffic-demand-matrix-based greedy route selection algorithm, B mis-associates the outcomes with its available actions (B has two available actions in the preceding example: choosing BD or BFD). This example is also an example of instability caused by route selection algorithms.

There is, however, a simple rational route selection algorithm that can choose the optimal route and maintain stability for B , if B does not restrict its route selection algorithm to always use the greedy strategy. This algorithm consists of an experimentation phase and a selection phase. At the beginning, B does not know the associated outcomes of choosing BD or BFD , thus it will first experiment with these two actions, one at a time. In this phase, B will fix its chosen action for enough amount of time, and observe the associated outcome of the chosen egress route (we assume that S will respond to B 's chosen egress route in bounded time). Using our notation in Section 3, B observes the set of network route selections $H = \{\{BD, SBD\}, \{BFD, SCD\}\}$. Denote $r = \{BD, SBD\}$ and $r' = \{BFD, SCD\}$. B then enters the selection phase. Since $u_B(r) > u_B(r')$, $A_B(r_{-B}) = A_B(r'_{-B}) = \{BD, BFD\}$, we have $U_B^\infty(\{r, r'\}) =$

²This example generalizes the oscillations of classical single-path adaptive routing where only latency is considered [2, 28].

$\{BD\}$. Therefore, B selects the optimal egress route BD , the only one in $U_B^\infty(\{r, r'\})$. Note that this simple algorithm conforms to the definition of rational route selection. On the other hand, the greedy algorithm does not since it chooses BFD infinitely often which is not in $U_B^\infty(\{r, r'\})$.

Therefore, depending on the route selection algorithms used by B , the example network may or may not experience routing instability, even if the local policies of the ASes in the network remain the same. This example thus serves as an example showing that, the stability of a network depends on the the route selection algorithms used by all ASes in the network.

5.3. Optimal and Stable Inbound-dependent Rational Route Selection by a Single AS

Generalizing the two-phase route selection algorithm for B in the example network in Section 5, Figure 4 specifies a rational route selection algorithm which can guarantee stability and optimality, when only AS i adopts this inbound-dependent route selection algorithm. Note that in Figure 4, r_i is a route selection constructed from the routes exported by AS i 's neighbors.

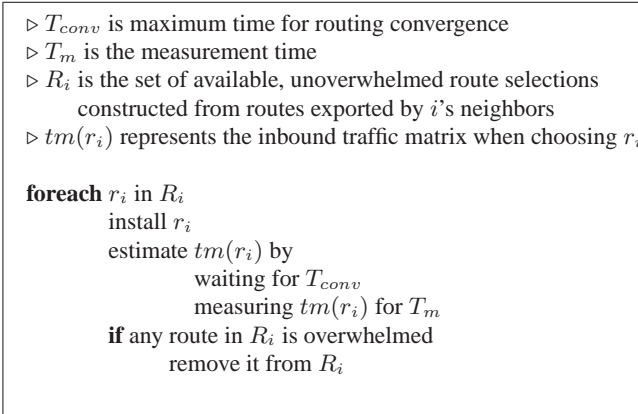


Figure 4. An inbound-dependent rational route selection algorithm by a single AS.

Specifically, in the context of Internet interdomain route selection, when ASes are constrained by Internet business considerations, Theorem 6 shows that the algorithm in Figure 4 can guarantee stability and optimality. Due to space limitation, we omit its proof, and note that an induction proof can be constructed.

Theorem 6 *The network converges, and an AS i converges to its optimal outcome, if the following conditions are satisfied:*

1. *there is no provider-customer loop in the network;*
2. *all ASes except i adopt the typical export policy;*

3. *each AS prefers customer routes over peer/provider routes;*
4. *AS i adopts the route selection algorithm in Figure 4, and no other AS uses any inbound-dependent route selection.*

6. Inbound-Dependent Route Selection: Instability of Networks under any Rational Route Selection Algorithms

Unfortunately, with inbound-dependency, there exist networks which have no stable route selection under any rational route selection algorithms; that is, we can arbitrarily assign route selection algorithm to each AS, so long each algorithm is a rational route selection algorithm, the network has no stable route selection.

In particular, Figure 5 is such an example network. Similar to the network in Figure 2, this network is constructed to satisfy all constraints imposed by AS business considerations; thus, if there were no inbound dependency, the network has a unique stable route selection [17]. Also similar to the network in Figure 2, this network does not appear to be a pathological case and can well happen in practice. Note that this network is a heterogeneous network, where the ranking of routes at S is inbound independent; while A and B are inbound dependent.




Figure 5. An example with instability. D is the only destination.

The instability of the example network in Figure 5 under any rational route selection scheme is established by the following result:

Theorem 7 *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 :*

$$\forall r'_i \in A_i(r_{-i}^*), u_i(r_i^*, r_{-i}^*) \geq u_i(r'_i, r_{-i}^*).$$

Proof: Since $\{r[t]\}_{t=0}^\infty$ converges to r^* , there exists t' such that $\forall t \geq t', r[t] = r^*$. Since the sequence is consistent with rational route selection, there exists $t'' > t'$, such that $\forall t > t''$ and $\forall i, r_i[t] \in U_i(\{r[s]|t' \leq s < t\})$. Notice that $\{r[s]|t' \leq s < t\} = \{r^*\}$, by definition of U_i , we have that

$$\forall r'_i \in A_i(r_{-i}^*), u_i(r_i^*, r_{-i}^*) \geq u_i(r'_i, r_{-i}^*).$$

■

An analysis of all of the possible network route selections of the example in Figure 5 shows that no network route selection satisfies the condition in Theorem 7. As a result, the network cannot converge to a stable route selection, under any rational route selection algorithm.

To further understand the example, consider the dynamics. When A and B choose AD and BFD . The outcome is SAD since S ranks SAD higher than $SBFD$. Then A has incentive to change from AD to AED since A ranks $\{S\}AED$ higher than $\{S\}AD$. However, 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.

7. Conclusions and Future Work

In this paper, we have proposed the notion of rational route selection algorithms, where inferior routes are iteratively eliminated. We derive a sufficient condition to check the stability of a heterogeneous network so long the route selection algorithm of each AS is rational in the context. Applying our general result, we analyze the stability of interdomain route selection where an AS's ranking on routes depends on inbound traffic. We have shown that the common scheme of choosing the best routes according to the traffic-demand matrix of the preceding period could lead to instability, when the inbound traffic depends on route selection. We have also shown that there exist networks where routing will be unstable under any rational route selection algorithms, even when the ASes strictly follow the constraints imposed by AS business considerations.

The unstable network shown in Section 6 is particularly troubling in that it does not appear to be a pathological case, and thus could happen in practice. When we encounter such an unstable network setting in practice, there is still no satisfactory solution. Fundamentally, to stabilize the network, tradeoff between local optimality and global stability must be made. Thus to design a stable route selection protocol, the ASes in a network must be willing to look into the future, form the right coalition, and sacrifice short-term benefits. Previous work such as route suppression (*e.g.*, [23]) and route dampening (*e.g.*, [34]) represents interesting potential directions. However, how to design interdomain routing protocols where the tradeoff between stability and local optimality is explicitly made in an incentive-compatible way is still a major remaining challenge.

Acknowledgments

We thank Jiang Chen, Ronny Dakdouk, Joan Feigenbaum, Eric Friedman, Arvind Krishnamurthy, and Jennifer Rexford for valuable comments and discussions. The

connection between inbound-dependent route selection algorithms and traffic-demand-matrix-based algorithms is pointed out by Tim Griffin. We are grateful to his help.

References

- [1] M. Afegan and J. Wroclawski. On the benefits and feasibility of incentive based routing infrastructure. In *Proceedings of ACM SIGCOMM '04 Workshop on Practice and Theory of Incentives and Game Theory in Networked Systems*, Portland, OR, Sept. 2004.
- [2] D. Bertsekas and R. Gallager. *Data Networks*. Prentice-Hall, Second Edition, 1992.
- [3] T. Bressoud, R. Rastogi, and M. Smith. Optimal configuration for BGP route selection. In *Proceedings of IEEE INFOCOM '03*, San Francisco, CA, Apr. 2003.
- [4] R. Dakdouk, H. Wang, H. Xie, and Y. R. Yang. Interdomain routing as social choice: A new perspective. Technical Report YALEU/DCS/TR-1331, Yale University, July 2005.
- [5] N. Feamster, J. Borkenhagen, and J. Rexford. Guidelines for interdomain traffic engineering. *ACM SIGCOMM Computer Communications Review*, Oct. 2003.
- [6] N. Feamster and J. Rexford. Network-wide BGP route prediction for traffic engineering. In *Proceedings of ITCOM*, Boston, MA, Aug. 2002.
- [7] J. Feigenbaum, D. Karger, V. Mirrokni, and R. Sami. Subjective-cost policy routing. Technical Report YALEU/DCS/TR-1302, Yale University, Sept. 2004.
- [8] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. In *Proceedings of the 21st ACM Symposium on Principles of Distributed Computing (PODC)*, pages 173–182, Monterey, CA, July 2002.
- [9] J. Feigenbaum, R. Sami, and S. Shenker. Mechanism design for policy routing. In *Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 11–20, St. John's, Newfoundland, Canada, July 2004.
- [10] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs. Locating Internet routing instabilities. In *Proceedings of ACM SIGCOMM '04*, Portland, OR, Aug. 2004.
- [11] A. Feldmann and J. Rexford. IP network configuration for intradomain traffic engineering. *IEEE Network Magazine*, pages 46–57, Sept./Oct. 2001.
- [12] B. Fortz, J. Rexford, and M. Thorup. Traffic engineering with traditional IP routing protocols. *IEEE Communication Magazine*, Oct. 2002.
- [13] E. Friedman. Asynchronous learning in decentralized environments: A game theoretic approach. In K. Tumer and D. Wolpert, editors, *Collectives and the Design of Complex Systems*. Springer-Verlag, 2004.
- [14] 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.

- [15] E. Friedman, M. Shor, S. Shenker, and B. Sopher. An experiment on learning with limited information: Nonconvergence, experimentation cascades, and the advantage of being slow. *Games and Economic Behavior*, 47(2):325–352, 2004.
- [16] L. Gao, T. G. Griffin, and J. Rexford. Inherently safe backup routing with BGP. In *Proceedings of IEEE INFOCOM '01*, Anchorage, AK, Apr. 2001.
- [17] L. Gao and J. Rexford. Stable Internet routing without global coordination. *IEEE/ACM Transactions on Networking*, 9(6):681–692, Dec. 2001.
- [18] R. Gao, C. Dovrolis, and E. W. Zegura. Interdomain ingress traffic engineering through optimized AS-path prepending. In *Proceedings of Networking'05*, Waterloo, Ontario, Canada, May 2005.
- [19] D. Goldenberg, L. Qiu, H. Xie, Y. R. Yang, and Y. Zhang. Optimizing cost and performance for multihoming. In *Proceedings of ACM SIGCOMM '04*, Portland, OR, Aug. 2004.
- [20] R. Govindan and A. Reddy. An analysis of Internet interdomain topology and route stability. In *Proceedings of IEEE INFOCOM '97*, Kobe, Japan, Apr. 1997.
- [21] T. G. Griffin, A. D. Jaggard, and V. Ramachandran. Design principles of policy languages for path vector protocols. In *Proceedings of ACM SIGCOMM '03*, Karlsruhe, Germany, Aug. 2003.
- [22] 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.
- [23] T. G. Griffin and G. Wilfong. A safe path vector protocol. In *Proceedings of IEEE INFOCOM '00*, Tel Aviv, Israel, Mar. 2000.
- [24] *On Stable Route Selection for Interdomain Traffic Engineering: Models, Analysis, and Guidelines*, Boston, MA, Oct. 2005.
- [25] A. Jaggard and V. Ramachandran. Robustness of class-based path-vector systems. In *Proceedings of the 12nd International Conference on Network Protocols (ICNP) '04*, Berlin, Germany, Oct. 2004.
- [26] A. Jaggard and V. Ramachandran. Relating two formal models of path-vector routing. In *Proceedings of IEEE INFOCOM '05*, Miami, FL, Apr. 2005.
- [27] R. Johari and J. N. Tsitsiklis. Routing and peering in a competitive Internet. Available at: <http://web.mit.edu/jnt/www/publ.html>, Jan. 2003.
- [28] A. Khanna and J. Zinky. The revised ARPANET routing metric. *ACM SIGCOMM Computer Communications Review*, Sept. 1989.
- [29] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed Internet routing convergence. In *Proceedings of ACM SIGCOMM '00*, Stockholm, Sweden, Aug. 2000.
- [30] C. Labovitz, G. R. Malan, and F. Jahanian. Internet routing instability. In *Proceedings of ACM SIGCOMM '97*, Cannes, France, Sept. 1997.
- [31] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Experiences applying game theory to system design. In *Proceedings of ACM SIGCOMM '04 Workshop on Practice and Theory of Incentives and Game Theory in Networked Systems*, Portland, OR, Sept. 2004.
- [32] R. Mahajan, D. Wetherall, and T. Anderson. Towards coordinated interdomain traffic engineering. In *Proceedings of Third Workshop on Hot Topics in Networks (HotNets-III)*, San Diego, CA, Nov. 2004.
- [33] R. Mahajan, D. Wetherall, and T. Anderson. Negotiation-based routing between neighboring domains. In *Proceedings of USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI '05)*, San Francisco, CA, May 2005.
- [34] Z. M. Mao, R. Govindan, G. Varghese, and R. H. Katz. Route flap damping exacerbates Internet routing convergence. *Computer Communication Review*, 32(4):221–233, 2002.
- [35] P. Milgrom and J. Roberts. Adaptive and sophisticated learning in normal form games. *Games and Economic Behaviors*, 3:82–100, 1991.
- [36] M. J. Osborne and A. Rubenstein. *A Course in Game Theory*. The MIT Press, 1994.
- [37] B. Quoitin, S. Uhlig, C. Pelsser, L. Swinnen, and O. Bonaventure. Interdomain traffic engineering with BGP. *IEEE Communications Magazine*, 41(5):122–128, May 2002.
- [38] J. L. Sobrinho. Network routing with path vector protocols: Theory and applications. In *Proceedings of ACM SIGCOMM '03*, Karlsruhe, Germany, Aug. 2003.
- [39] L. Subramanian, M. Caesar, C. T. Ee, M. Handley, Z. M. Mao, S. Shenker, and I. Stoica. Towards a next generation inter-domain routing protocol. In *Proceedings of Third Workshop on Hot Topics in Networks (HotNets-III)*, San Diego, CA, Nov. 2004.
- [40] R. Teixeira, T. Griffin, A. Shaikh, and G. Voelker. Network sensitivity to hot-potato disruptions. In *Proceedings of ACM SIGCOMM '04*, Portland, OR, Aug. 2004.
- [41] R. Teixeira, A. Shaikh, T. Griffin, and J. Rexford. Dynamics of hot-potato routing in IP networks. In *Proceedings of Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, New York, NY, June 2004.
- [42] K. Varadhan, R. Govindan, and D. Estrin. Persistent route oscillations in inter-domain routing. *Computer Networks*, 32(1):1–16, 2000.
- [43] H. Wang, R. K. Chang, D.-M. Chiu, and J. C. Lui. Characterizing the performance and stability issues of the AS path prepending method: Taxonomy, measurement study and analysis. In *Proceedings of ACM SIGCOMM Asia Workshop*, Beijing, China, Apr. 2005.
- [44] H. Wang, H. Xie, Y. R. Yang, L. E. Li, Y. Liu, and A. Silberschatz. On stable route selection for interdomain traffic engineering: Models, analysis, and guidelines. Technical Report YALEU/DCS/TR-1316, Yale University, Feb. 2005.