

# Stable iBGP through Selective Path Dissemination

Ravi Musunuri                      Jorge A. Cobb  
Department of Computer Science (EC 31)  
The University of Texas at Dallas  
Richardson, TX 75083-0688  
{musunuri, cobb}@utdallas.edu

## ABSTRACT

In Internet, routing between Autonomous Systems (ASes) is performed by Border Gateway Protocol (BGP). Neighboring routers in different ASes share their routing information by using external BGP (eBGP), while routers in the same AS share their routing information about external destinations by using internal BGP (iBGP). iBGP employs route reflection clustering to solve the scalability problems due to number of iBGP peering sessions required. This clustering, however can cause anomalies such as routing loops and a failure to converge to a stable set of paths.

In this paper, we present a Selective Path Dissemination (SPD) protocol, which avoids iBGP anomalies specific to clustering, and we also prove the correctness of the SPD protocol. Our solution requires minor changes to route reflection clustering method. In our proposed method, each border router would act as the reflector and each interior router is clustered with its nearest border router. Furthermore, each border router selectively disseminates paths to each of interior routers in border router's cluster.

## KEY WORDS

iBGP, Route Reflection Clustering, Divergence, Routing Loops

## 1 Introduction

At its highest level, the Internet is a collection of interconnected Autonomous Systems (ASes), where each AS is a collection of interconnected networks under the management of a single organization. Routing of IP datagrams within an AS is performed at the discretion of the organization. Therefore, each AS is free to choose any routing algorithm for routing within the AS (intra-AS routing). Although the intra-AS routing protocol is chosen by each AS, routing between ASes (inter-AS routing) must be performed by a common protocol executed by all ASes. Because ASes are interconnected, the Internet may be viewed as an AS-graph, i.e., a graph where each node is an AS. Each router in an AS stores in its memory the paths along the AS-graph that router has chosen to reach every other AS. In order to learn new paths and to propagate existing paths, each router informs its neighbor of the paths the router has chosen. The purpose of exchanging an

entire path to each destination as opposed to exchanging a distance vector is to avoid long-lived routing loops. The propagation of these paths from one AS to another is performed through the Border Gateway Protocol (BGP). Neighboring routers in different ASes share their chosen paths by using the *external* Border Gateway Protocol (eBGP), while routers in the same AS share their chosen paths to external destinations by using the *internal* Border Gateway Protocol (iBGP).

Inside the AS, each border router shares the paths, it has chosen to external destinations, to every router. This sharing of paths requires Transmission Control Protocol (TCP) connection, also known as *iBGP peering session*, from each border router to all the interior routers in the AS. For scalability, routers within the AS are divided into several disjoint *clusters*. Each cluster will have special router known as the *reflector*. Each border router only shares chosen paths with reflector in the border router's cluster. Each reflector is responsible for sharing chosen paths to every reflector inside its AS and other routers inside its cluster. Multi Exit Discriminator (MED) attribute is used to discriminate the links between same ASes.

Both eBGP and iBGP have been plagued with a multitude of anomalies. eBGP has the divergence anomaly [1, 2, 3], in which routers would exchange path update messages continuously without settling for stable set of paths. One reason for divergence is, each AS may rank each path according to some local policy, that is, paths are ranked independently at each AS. Therefore, the ranking of paths at different ASes may conflict with each other, leading to unstable behavior. There has been many solutions [2, 4, 5, 6, 7] proposed to solve the divergence problem in eBGP but none are satisfactory. iBGP also has divergence and routing loops [8, 9, 10, 11, 12] anomalies. Griffin et al. [9] observed that this iBGP anomalies could occur due to clustering. In particular, these anomalies may occur if interior routers are not using paths through their nearest border router with a path to reach a particular destination. No solution has been proposed in the current literature to solve the iBGP anomalies specific to the clustering. In this paper, we are proposing a solution to solve the anomalies specific to iBGP clustering.

We present a Selective Path Dissemination (SPD) protocol, which avoids iBGP anomalies specific to clustering and we also prove the correctness of SPD protocol. Our solution requires minor changes to existing clustering method. In our method, each border router acts as the reflector and each interior router is clustered with its closest border router. Furthermore, each border router selectively disseminates paths to each of interior routers in border router’s cluster.

This paper uses the following *conventions* in the figures. We refer to a link as *internal* link of an AS, if it is connecting two routers in the same AS. Otherwise, we refer that link as *external* link. All routers shown in the figures are trying to find a best path to special destination AS  $d$ . ASes are shown as ellipses, border routers are named as  $B_i$ , interior routers are named as  $I_i$ , and external paths to destination AS  $d$  through border router  $B_i$  are named  $P_i$ . Clusters are represented as  $C_i$  and reflector is shown in bold face letter in each cluster. We also assume external paths  $P_i$  have the same local preference values and AS Path attribute lengths. The weights on the links represent Interior Gateway Protocol (IGP) metrics. Solid lines indicate physical links between routers and dashed lines indicate iBGP peering sessions between routers. We are also using the following *notation* in this paper. We refer border routers in  $AS$ , whose current path to  $d$  is passing through only external links as the *feasible* border routers. Otherwise we refer them as *infeasible* border routers. Interior router  $I_i$ ’s *nearest* border router is referred to as  $\beta(I_i)$  and *nearest feasible* border router is referred to as  $\psi(I_i)$ .

The remaining sections are organized as follows. In Section 2, we give a brief overview of BGP, and route reflection scaling technique used in iBGP. In Section 3, we discuss the related work that addresses the instability problems in iBGP. In Section 4, we will explain different iBGP anomalies with examples. In Section 5, we present SPD protocol and prove the protocol correctness. In Section 6, we give our concluding remarks and areas of future work.

## 2 BGP Overview

BGP [14] is de-facto standard for sharing of the routing information between ASes. Neighbors in different ASes are known as eBGP peers and TCP connection between them is known as an eBGP peering session. Any two routers, who belong to the same AS are called as iBGP peers and TCP connection between them is called as an iBGP peering session, even if they are not physically connected by single link. In BGP, the update message is used by the routers to advertise the chosen path and to withdraw previously advertised infeasible paths. Update message carries *origin*, *next hop*, *AS path*, *MED* path attributes. Origin attribute contains the origin AS identifier of the update message. The Next Hop attribute, which contains the IP

address of next hop router, is used to forward IP datagrams to next-hop router along the current chosen path. AS Path attribute stores list of ASes along the path, which is used to avoid loops. MED is used to inform which link is preferred over the other, when there are multiple links between same ASes. Smaller MED value link is preferred over higher MED value link.

### 2.1 Route Selection Process

In this section we will explain the commonly used route selection process at the BGP routers. BGP routers select the best path from the available paths set by using the following steps. Route selection process proceeds to the next step if there are multiple paths in the previous step and stops if there is only one path.

1. Router selects the path with the best local preference value.
2. If there are multiple such paths then router chooses the path with minimum AS Path attribute length.
3. If there are multiple such paths then router chooses the path through a BGP peer with least MED for each of the neighboring ASes.
4. If there are multiple such paths then one of following may determine the best path.
  - If the router has path through atleast one eBGP peer then the router chooses the path through an eBGP peer with least IGP metric between itself and eBGP peer.
  - If router has no path through eBGP peers then the router chooses the least IGP metric path to the border router, which has a path to destination AS.
5. Finally if there are multiple paths, use some deterministic tie breaker, such as least Next Hop path attribute to find the best path.

### 2.2 iBGP Explained

BGP peers, within the same AS, use iBGP protocol to exchange paths they have chosen for the destinations outside AS. The primary difference between iBGP and eBGP is the mechanism by which router checks for the loops in the paths. As mentioned earlier, eBGP uses AS Path attribute to avoid loops. In iBGP, all routers belong to same AS, iBGP can’t use AS Path attribute to avoid loops. Interior routers avoid loops by not forwarding chosen paths to any other iBGP peers. Due to which, each border router needs to maintain an iBGP peering session with every other router. But this full-mesh peering method is not scalable. Route reflection [15], Confederation [12] are two mechanisms that can alleviate the scaling problem. In this paper we are only considering iBGP anomalies due to route reflection clustering.

### 2.3 Route Reflection Clustering

In route reflection clustering, all routers inside the same AS are divided into disjoint sets known as *clusters*. Each cluster will have special router known as the *reflector* and other routers in the cluster are known as *clients*. All the reflectors inside the same AS will maintain a full-mesh of iBGP peering sessions with each other and each reflector also maintains an iBGP peering session with every client in reflector's cluster. This clustering significantly reduces the iBGP peering sessions required. In fig.1, AS is divided into  $\{I_1, B_1\}, \{I_2, B_2\}$  clusters with reflectors  $B_1$  and  $B_2$  respectively. Reflectors  $B_1, B_2$  maintain full-mesh iBGP peering sessions and each reflector also maintains iBGP peering session with clients in it's cluster.

In this clustering, border routers are required to forward chosen paths only to cluster's reflector. Reflectors are responsible for forwarding chosen paths to other reflectors and every cluster's client router. Reflector advertises chosen paths in the following way.

- If the reflector received a path update message from one of its clients then the reflector advertises the chosen path to all of other clients inside its cluster and to every other reflector.
- If the reflector received a path update message from an eBGP peer then the reflector advertises the chosen path to all the clients inside its cluster and to every other reflector.
- If the reflector received a path update message from other reflector then the reflector advertises the chosen path to every client in its cluster.

### 3 Related Work

Anomalies specific to the clustering can occur even with stable eBGP and without even considering the MED values. In [9], authors have studied the correctness of iBGP and mentioned sufficient conditions to avoid anomalies specific to the clustering. There are also many studies related to the divergence problems if both MED and clustering are considered [13, 8, 10, 11, 12]. There are few solutions proposed to solve the divergence problem due to MED and clustering [11, 13]. The important idea in these solutions is that each router, in addition to its best path, also advertises some additional paths. These solutions are not scalable. In this paper, we are proposing a new solution to solve the iBGP anomalies specific to clustering. We do not consider the MED attribute values in our solution.

### 4 iBGP Anomalies due to Clustering

In this section we explain the forwarding anomaly and the divergence anomaly in iBGP with examples. These anomalies occur due to improper clustering of routers and incorrect placement of reflectors.




Figure 1. Forwarding anomaly due to clustering.

#### 4.1 Forwarding Anomaly

Lets us consider the interior Router  $I_i$ , whose current path to destination  $d$  is passing through the border Router  $B_k$ . Also assume, shortest path between  $I_i$  and  $B_K$  is passing through an intermediate Router  $I_j$ . But Router  $I_j$ 's current path to same destination  $d$  is passing through a different border router  $B_m$ . IP datagrams from router  $I_i$  get deflected at router  $I_j$ , which is known as the forwarding anomaly. Forwarding anomalies complicate the protocol debugging. In the worst case, these deflections could lead to routing loops. Next, we will discuss an example with routing loop.

Figure 1 shows an AS with routing loop, which is taken from [9]. In this example, AS is divided into two clusters  $\{B_1, I_1\}$  and  $\{B_2, I_2\}$ .  $B_i (i = 1, 2)$  acts as the reflector in each cluster. Router  $B_i$  always chooses  $P_i$  as the current path by using the route selection process mentioned in Section 2.1. Notice that, the route selection process favors the external links over the internal links at the border routers. Client  $I_i$  only learns about the path  $P_i$  from its reflector  $B_i$ .  $I_1$ 's current path is passing through the border router  $B_1$  and the shortest path between  $I_1$  and  $B_1$  is  $I_1 \rightarrow I_2 \rightarrow B_1$ . But  $I_2$ 's path is passing through a different border router  $B_2$ . IP datagrams from  $I_1$  to the destination  $d$  get deflected at  $I_2$  and IP datagrams from  $I_2$  to the destination  $d$  get deflected at  $I_1$ . Hence, there is a forwarding loop between  $I_1$  and  $I_2$  routers. Next, we will discuss the reasons for this anomaly and how we can avoid loops.

Each interior router  $I_i$  should choose the path through  $\psi(I_i)$ . But, Router  $I_1$ 's path is passing through  $B_1$ , which is not  $\psi(I_1)$ . Similarly  $I_2$ 's path is passing through the border router  $B_2$ , which is not  $\psi(I_2)$ . Therefore, We can avoid routing loops by making border routers as the reflectors and each interior router  $I_i$  is clustered with  $\beta(I_i)$ . In this example, each interior router chooses the path through  $\psi(I_i)$ , which is equal to  $\beta(I_i)$ . The anomaly free clustering is shown in fig. 2.

#### 4.2 Divergence Anomaly

In the divergence anomaly, routers would exchange the path update messages continuously. Router's current paths




Figure 2. Forward anomaly free clustering.




Figure 3. Divergence anomaly due to clustering.

would change continuously without settling for the stable set of paths. iBGP can also diverge due to combination of clustering and MED attribute [13] [10]. But, the divergence anomaly we are considering is specific to the clustering of routers. Next, we will explain the divergence anomaly with an example.

Divergence example taken from [16] is shown in the fig.3. In this example, each  $I_i$  acts as the route reflector for the cluster  $\{B_i, I_i\}$ .  $B_i$  always chooses  $P_i$  as the current path by using the route selection process mentioned in Section 2.1. Router  $I_i$  always prefers  $P_{i+1}$ <sup>1</sup> over  $P_i$ . Following explains the divergence problem at  $I_i$  routers.

1. Lets assume  $I_1$ 's current path is  $P_2$ ,  $I_2$ 's current path is  $P_2$  and  $I_3$ 's current path is  $P_3$ .
2. Next, If  $I_2$  receives the path update message from  $I_3$  then  $I_2$  withdraws  $P_2$  and changes its current path to  $P_3$ . So  $I_1$  changes its current path to  $P_1$ .
3. Next, If  $I_3$  receives the path update message from  $I_1$  then  $I_3$  withdraws  $P_3$  and changes its current path to  $P_1$ . So  $I_2$  changes its current path to  $P_2$ .
4. Next, If  $I_1$  receives update from  $I_2$  then  $I_1$  withdraws  $P_1$  and changes its current path to  $P_2$ . So  $I_3$  changes its current path to  $P_3$ .

$I_i$  routers would continuously exchange path update messages in the cyclic manner as above. They will never agree on stable set of paths. Next, we will discuss reasons for this anomaly and how we can avoid it.




Figure 4. Divergence anomaly free clustering.

Border routers will have better knowledge about announcing new feasible paths and withdrawing infeasible paths. In step-2 of divergence anomaly example,  $I_2$  withdraws the path  $P_2$ . Due to which,  $I_1$  changes its current path to  $P_1$ . But  $P_2$  could still be a feasible path.  $I_2$  is unnecessarily withdrawing the path  $P_2$ . Border router  $B_2$  has better knowledge about whether  $P_2$  is a feasible path or not? Therefore, border routers should act as the reflectors. We can avoid the divergence anomaly by making  $B_i$  as the reflector and each interior router  $I_i$  is clustered with  $\beta(I_i)$  as shown in fig.4. Every interior router  $I_i$  chooses the path through  $\psi(I_i)$ .

## 5 Selective Path Dissemination (SPD) Protocol

In this section we will explain selective path dissemination protocol, which avoids iBGP anomalies specific to the clustering. From above discussion in examples, it seems like we can solve iBGP anomalies by making two minor changes to the clustering mechanism. First, by making border routers as the reflectors. Second, each interior router is clustered with the nearest border router. But these changes may not solve iBGP anomalies in all situations. In particular, if  $I_i$ 's nearest border router  $\beta(I_i)$  is different from nearest feasible border router  $\psi(I_i)$ .

Let us consider an example shown in fig.5, to motivate extra changes required to our protocol. In this example, each of interior router  $I_i$  ( $i = 1, 2$ ) is clustered with its  $\beta(I_i)$ . Other border routers  $B_3$  and  $B_4$  form separate clusters. Border routers  $B_1$  and  $B_2$  are infeasible border routers. Border routers  $B_1$  and  $B_2$  chose their paths through their nearest border routers  $B_3$  and  $B_4$  respectively.  $B_1$  advertises path  $P_3$  to  $I_1$  and  $I_1$ 's shortest path to  $B_3$  is  $I_1 \rightarrow I_2 \rightarrow B_3$ . By similar arguments,  $I_2$  chooses path  $I_2 \rightarrow I_1 \rightarrow B_4$ . IP datagrams from  $I_1$  get deflected at  $I_2$  and IP datagrams from  $I_2$  get deflected at  $I_1$ . Hence, there is routing loop between routers  $I_1$  and  $I_2$ . We can avoid this loop by using selective path dissemination at border routers.

In our method, border routers act as the reflectors. So all border routers will have full knowledge about available

<sup>1</sup> mod 3 is applied to subscripts in this example description




Figure 5. Protocol Motivation Example.

paths through all other border routers. Infeasible border routers, instead of sending chosen path to all interior routers in their cluster, they should selectively advertise paths through  $\psi(I_i)$  to each interior router  $I_i$ . Each interior router  $I_i$  will receive an update message with path through nearest feasible border router  $\psi(I_i)$  from nearest border router  $\beta(I_i)$ .

This extra change will solve iBGP anomalies completely. But how does border routers know about each interior router  $I_i$ 's  $\psi(I_i)$ ? If AS is using link state routing, such as Open Shortest Path First (OSPF), for intra-AS routing then every router will have knowledge about the full topology. If AS is using distance vector protocol, such as Routing Information Protocol (RIP), for intra-AS routing then this assumption is not true. To make SPD protocol independent from intra-AS routing protocol, infeasible border router sends a *request()* message to each interior router in its cluster. Interior router sends back a *reply()* message with the sorted list of border router identifiers.

### 5.1 Pseudo-code at Interior Router

Fig. 6 shows pseudo-code for SPD protocol at the interior router (*ir*). In this section we will present pseudo-code for SPD protocol at interior router (*ir*). Each interior router maintains a sorted list of border router identifiers (*brs*) and current path (*cp*) variables. This list contains routers identifiers from the nearest to the farthest border router. Routers will use deterministic tie-breaker, such as smallest border router identifier, to sort this border routers list. Interior router has two events. In first event, if *ir* receives an *update(cp)* message from  $\beta(ir)$  then *ir* updates the current path (*cp*) variable. In second event, If *ir* receives a *request()* message from  $\beta(ir)$  then *ir* sends back a *reply(brs)* message.

### 5.2 Pseudo-code at Border Router

Fig. 7 shows pseudo-code for SPD protocol at the border router (*br*). Each border router maintains the current path (*cp*) variable. Each border router (*br*) has three events. In first event, if *br* receives *update(p)* message from another border router then *br* does the following actions. *br* updates its current path (*cp*) by calling *best\_path* function. Boolean expression (*border\_router(cp) = br*)

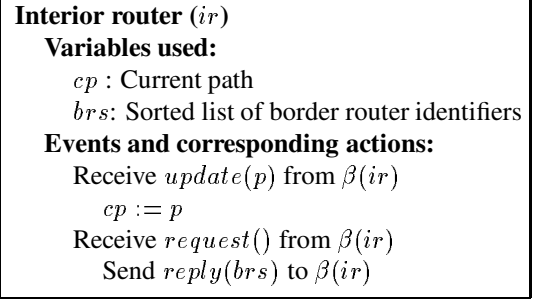


Figure 6. SPD Protocol for Interior Router

is true if *br* router is feasible. If *br* is feasible then *br* forwards the chosen path to all interior routers in its cluster. Otherwise *br* will send a *request()* message to all interior routers in its cluster. In second event, If *br* receives an *update* message from eBGP peer then *br* does the following actions. *br* updates the current path (*cp*) by calling *best\_path* function and If *br* is feasible then *br* forwards the chosen path to all interior routers in its cluster. In third event, If *br* receives a *reply(list br)* message from interior router *ir* then *br* does the following actions. *br* selectively finds the best path through  $\psi(ir)$  by using function *selective\_best\_path(brs, ir)* function and forwards this path to *ir*.

### 5.3 Correctness Proof

In this section we will formally prove the correctness of SPD protocol

**Lemma 1** Every Router *R* always chooses its path through nearest feasible border router  $\psi(R)$ .

**Proof:** In SPD protocol, border router  $B_i$ s are acting as the reflectors. Every border router will have a full knowledge of available paths to the AS. Hence, every border router  $B_i$  always chooses the path through  $\psi(B_i)$ . Every interior router  $I_i$  is clustered with the nearest border router  $\beta(I_i)$ . Now we need to consider two types of clusters. In first case, If the cluster's border router is feasible then every interior router selects path through  $\psi(I_i)$ , which is equal to  $\beta(I_i)$ . In second case, If the cluster's border router is not feasible then every interior router sends a sorted list border router identifiers to border router. Border router selectively advertises the path through  $I_i$ 's nearest feasible border router  $\psi(I_i)$ . Which proves the Lemma.  $\square$

**Lemma 2** If any interior router  $I_i$  chooses path through border router  $B_i$  then all the routers along the shortest path between  $I_i$  and  $B_i$ , including  $B_i$ , should choose path through same border router  $B_i$ .

**Proof:** We prove this by contradiction using fig. 8. Lets assume Interior router  $I_i$  chooses path through border router  $B_i$ . From Lemma 2,  $\psi(I_i) = B_i$ . At some intermediate router  $I_j$  along the shortest path between  $I_i$  and  $B_i$  chooses path through different border router  $B_j$ . Therefore,  $B_j$  should be nearer to  $I_i$  than  $B_i$  and  $\psi(I_i) = B_j$ . Which contradicts our assumption.  $\square$

### Border router ( $br$ )

#### Variables used:

$cp$  : Current path.

#### Events and corresponding actions:

Receive  $update(p)$  from iBGP peer

$cp := best\_path(update(p))$

If ( $border\_router(cp) = br$ )

Send  $update(cp)$  to interior routers

Else

Send  $request()$  to interior routers

Receive  $update(p)$  from eBGP peer

$cp := best\_path(update(p))$

If ( $border\_router(cp) = br$ )

Send  $update(cp)$  to interior routers

Receive  $reply(br)$  from interior router  $ir$

$\beta(ir) = selective\_best\_path(br)$

Send path  $update(\beta(ir))$  to  $ir$

#### Functions used:

- $best\_path(update(p))$ : Updates the available paths database with information from new  $update$  message and returns the new best path by using the route selection process.
- $border\_router(path)$ : Returns border router's identifier along the path.
- $selective\_best\_path(list, ir)$ : Returns  $path$  through the router  $\psi(ir)$  from the available set paths.

Figure 7. SPD Protocol for Border Router




Figure 8. Proof of Lemma 2.

**Theorem 1** *SPD protocol solves forwarding anomaly.*

**Proof:** From Lemma 2, If interior router  $I_i$  chooses path through  $B_i$  then all routers along the shortest path between  $I_i$  and  $B_i$  also choose path through  $B_i$ . IP datagrams from  $I_i$  would never get deflected.  $\square$

**Theorem 2** *SPD protocol solves divergence anomaly.*

**Proof** In SPD protocol border routers are acting as reflectors. Every border router will have full knowledge of available paths. Every border router chooses its path from the same set of available paths. Therefore, paths at border routers would never diverge. Furthermore, every interior router  $I_i$  chooses path through  $\psi(I_i)$ . Therefore, paths at interior routers would never diverge.  $\square$

## 6 Conclusions and Future work

BGP has been plagued with multitude of anomalies. Many solutions have been proposed for the divergence problem in eBGP. In this paper, we presented Selective Path Dissemination protocol, which solves iBGP anomalies specific to clustering with minor changes to clustering method. We also proved the correctness of SPD protocol. Our solution also solves MED based anomalies in iBGP. Due to space limitations we are unable address MED based anomalies in this paper. Our solution requires border routers to act as the reflectors. In future, we want to remove this restriction and allow any router to act as the route reflector. We also want to investigate how our protocol adjusts to node failures and link failures.

## References

- [1] Varadhan, K., Govindan, R., Estrin, D., Persistent Route Oscillations in Inter-Domain Routing, *Computer Networks*, Amsterdam, Netherlands.
- [2] Griffin, T.G., Shepherd, F.B., Wilfong, G., The Stable Paths Problem and Interdomain Routing, *IEEE/ACM Transactions on Networking*, Vol. 10, Issue 2, Apr. 2002, Pages 232-243.
- [3] Ratul Mahajan, David Wetherall, Tom Anderson, Understanding BGP Misconfiguration, *Proceedings of the ACM SIGCOMM Conference*, PA, Aug. 19-23, 2002.
- [4] Villamizar, C., Chandra, R., Govindan, R., *IETF RFC-2439: BGP route flap damping*, Nov., 1998.
- [5] Ramesh G., Alaettinoglu, C., George Eddy, David Kessens, Satish Kumar, WeeSan Lee, An Architecture for Stable, Analyzable Internet Routing, *IEEE Network Magazine*, Jan-Feb 1999.
- [6] Cobb, J.A., Gouda, M.G., Musunuri, R., *A Stabilizing Solution to the Stable Path Problem. Self-Stabilizing Systems, Self-Stabilizing Systems 2003*, Pages 169-183.
- [7] Lixin Gao, Jennifer Rexford, Stable internet routing without global coordination, *IEEE/ACM Transactions on Networking*, Vol. 9, Issue 6, Dec. 2001, Pages 681-692.
- [8] McPherson, D., Gill, V., Walton, D., Retana, A., *IETF RFC-3345: Border Gateway Protocol (BGP) Persistent Route Oscillation Condition*, Aug. 2002.
- [9] Griffin, T.G., Wilfong, G., On the Correctness of iBGP Configuration, *Proc. of ACM SIGCOMM Conference*, Pittsburgh, PA, Aug., 19-23, 2002.
- [10] Griffin, T.G., Wilfong, G., Analysis of the MED Oscillation Problem in BGP, *Proc. of ICNP Conference*, Paris, France, Nov., 12-15, 2002.
- [11] Walton, D., Cook, D., Retana, A., and Scudder, J., *BGP persistent Route Oscillation Solution*, IETF Internet draft, draft-walton-bgp-route-oscillation-stop-00.txt, Work in progress, May 2002.
- [12] Traina, P., McPherson, D., Scudder, J., *IETF RFC-3065: Autonomous System Confederations for BGP*, Feb., 2001.
- [13] Anindya B., Chih-Hao Luke Ong, April Rasala, Shepherd, F.B., and Wilfong, G., Route Oscillations in I-BGP with Route Reflection, *Proc. of ACM SIGCOMM Conference*, Pittsburgh, PA, Aug. 19-23, 2002.
- [14] Rekhter, Y. and Li, T., *IETF RFC-1771: A Border Gateway Protocol 4 (BGP-4)*, Mar. 1995.
- [15] Bates, T., Chandra, R., *IETF RFC-1996: BGP Route Reflection - An Alternative to Full Mesh iBGP*, June, 1996.
- [16] Scudder, J., G., and Dube, R., BGP Scaling Techniques Revisited, *ACM Computer Communication Review*, vol. 29, no. 3, Oct. 1999.