Various Approaches used in balancing Inbound Traffic Engineering for Multihomed ASs.

Lata L. Ragha

Abstract— Many Internet service Providers tune the configuration of the Boarder Gateway Protocol on their routers to control their Traffic. Controlling inbound traffic is quite difficult than outbound traffic over multiple links. In this paper five different approaches: modifying the MED attribute, selective announcement, prefix splitting, AS PATH prepending and BGP communities; used for balancing inbound traffic for multhomed ASs are discussed and summarized. The readers of this paper are intended to have the basic knowledge of routing and the Boarder Gateway Protocol (BGP).

***Index Terms*—AS, ASPP, BGP, Interdomain, Intradomain, MED, multihomed AS, stub AS, transit AS,**

I. Introduction

Intradomain and *interdomain* are the two levels of today's Internet routing architecture. The former refers to routing within a domain or an autonomous system (AS), while the latter refers to the routing between ASs. An AS is defined as "a connected group of one or more IP prefixes run by one or more network operators which has a single and clearly defined routing policy" [1], and each AS is uniquely identified by an AS number. An IP prefix is the network part of an IP address that is examined by routers to make forwarding decisions. In general, there are two types of AS, namely, *transit AS* and *stub AS*. A transit AS provides Internet connectivity to other ASes by forwarding all types of traffic across its network. A stub AS, on the other hand, does not provide transit service for other ASes and only sends or receives its own traffic. Figure 1 shows an example of interconnected ASs. Both AS1 and AS9 are *stub ASs*, while AS2–AS8 are *transit ASs*. The interconnection of ASes can also be described by a *business relationship*. Major business relationships include the *provider-to-customer* relationship and the *peer-to-peer* relationship. These business relationships play a crucial role in shaping the structure of the Internet and the end-to-end performance characteristics [2]. From the viewpoint of AS relationship, stub ASes are those which have no customer (or client AS), while transit ASes are those with customers. Transit ASes without provider are called "tier-1" ASes.

ASes that have more than one provider are called *multihomed ASes*. Motivated by the need to improve network resilience and performance, there is an increasing number of enterprise and campus networks connecting to the Internet via multiple providers. These multihomed ASes, therefore, must undertake the task of engineering the traffic flowing in and out of the network through these multiple links. In Figure 1 each stub AS is multihomed to two transit ASs; thus, they can receive and send packets via both links at the same time. The figure also shows the end-to-end routing path from a host in AS9 to another host in AS1. The entire routing path is therefore composed of intradomain routing paths and interdomain routing paths, alternating between them.

Traffic engineering is another important problem to tackle at the routing layer. The general problem at hand is how to influence the traffic flowing into (inbound) and out of (outbound) an AS, such that a given set of performance objectives can be achieved. Traditionally, the traffic engineering problem does not concern stub ASs, because most of them are singlehomed. However, as the number of multihomed stub ASs has been increasing rapidly for the last few years, the problem of engineering the inbound and outbound traffic becomes very important for a large number of ASs in the Internet.

Using different inter-AS traffic engineering approaches, ASes can distribute traffic to satisfy their performance or cost constraints [4]. The focus of this paper is on the *inter-AS inbound traffic* balancing, which is known to be more difficult than the outbound traffic engineering problem because an AS generally cannot control the routing path for the inbound traffic.

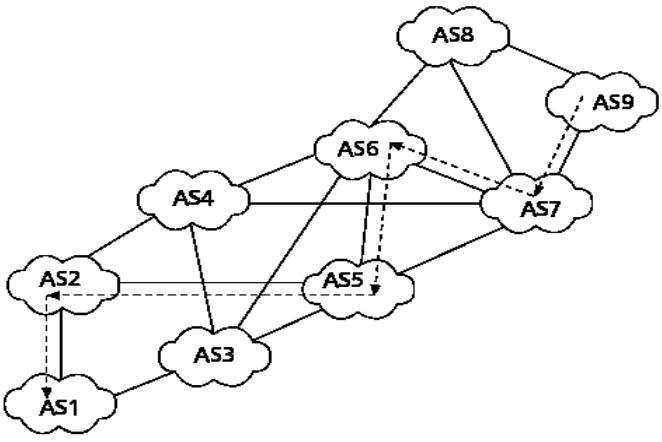

Figure 1. Interconnected Ass and forwarding path from stub AS9 to stub AS1

Lata L. Ragha is with the Department of Computer Engineering, Ramrao Adik Institute of Technology (affiliated to Mumbai University), Navi-Mumbai - 400706, India (e-mail: lata.ragha@gmail.com, lata.ragha@rediffmail.com).

II. BOARDER GATEWAY PROTOCOL

There are many routing protocols available for the intradomain level, such as Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Intermediate System to Intermediate System (IS-IS), the Border Gateway Protocol (BGP4) is the only standard for exchanging reachability information on the interdomain level [4]. BGP supports classless interdomain routing, and an important function of BGP is to facilitate *policy routing*. That is, each AS exercises its own preference for which routes to accept and where to further advertise them. To support such autonomous route decisions, a prefix announced in a BGP route advertisement is usually attached with a number of *path attributes*. A BGP router makes a route decision based on the values of the BGP AS-PATH attributes. For instance, AS-PATH contains the ASes through which the announcement for the prefix has passed. As an announcement is passed between ASes, each AS adds its AS number (ASN) to the AS-PATH attribute. This, by itself, is useful for the operators of the ASes to learn all the information of this route. Therefore, the final end-to-end forwarding path is essentially a result of the autonomous route decisions of the ASs between the two endpoints. Each link in Figure 1 may represent a BGP connection between two BGP routers in the respective ASs. Assume that AS1 advertises a prefix to the two links, which is in turn advertised to all BGP connections in the figure. Based on the final forwarding path, AS5's preferred next hop for the prefix is AS2 (instead of AS3), whereas AS6's preferred next hop is AS5 (instead of AS4 or AS3), and so on.

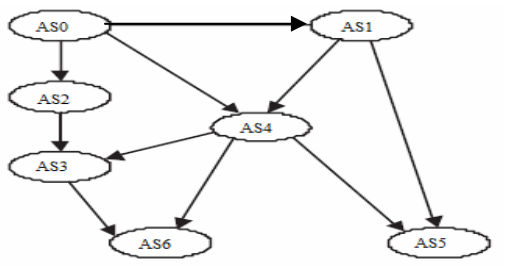

Figure 2. An example of network with different relationships

Different ASes have different business concerns, so there are different business agreements between ASes. BGP provides a mechanism to enforce business agreements made between two or more parties. This can be illustrated by the following example. In Figure 2, AS_3 and AS_4 are providers of AS_6, which implies AS_6 pays for the traffic going through the link $AS_3 - AS_6$ and the link $AS_4 - AS_6$. Imagine AS_3 wants to send traffic to AS_4. AS_6, being a customer to both AS_3 and AS_4, obviously does not want to provide transit service for its providers. To achieve this goal, AS_6 will not announce the reachability information of AS_3 (AS_4) to AS_4 (AS_3). In short, it is the role of an ISP's routing policy to enforce these kinds of business agreements.

BGP has two kinds of routing policies: *import routing policy* and *export routing policy* (also referred to as *import filtering* and *export filtering*). Import policy determines which routes should be accepted from a neighbor and the preference with which those routes should be treated, while export policy determines which routes should be advertised to a neighbor. If an AS accepts a route from a neighbor, it means this AS agrees to provide transit service for the traffic destined to the prefix of this route. If an AS advertises a route to one of its neighbours, it means this AS would like to accept traffic destined to the prefix of this route from this neighbour. Thus this kind of *routes filtering* is important and necessary for BGP to control how an ISP network is used by its neighbours.

BGP is a *policy-based* path vector routing protocol. In [5], the authors illustrate the popular policies adopted by ASes in the Internet are: (a) the *typical local preference import policy* and (b) the *selective announcement export policy*. Under the typical local preference policy, an AS prefers to use a customer link than a peering link to forward a packet, and it prefers to use a peering link than a provider link to forward a packet, provided that these links can reach the destination AS. This is natural since an AS does not need to pay for the traffic going through its customer link, while it must pay for the traffic going through its provider link. Under the selective announcement export policy, an AS would not announce the routes learned from its providers or peers to other providers and peers, thus an AS does not provide transit service between its providers or its peers. To illustrate, let us assume all ASes in Figure 2 obey "local preference" and "selective announcement" policies. Then routes with AS path (AS_5,AS_4,AS_6) or (AS_5,AS_1,AS_4,AS_3,AS_6) are considered legal or valid routes, while routes with AS path (AS_1,AS_0,AS_4,AS_6) would not appear in this network because AS_1 would select AS_4, instead of AS0 as the next hop to reach AS_6 according to the typical local preference. Also, route with AS path (AS_1,AS_4,AS_0) would not appear since AS_4 would not announce AS path (AS_4,AS_0) to AS_1 according to the "*selective export policy*".

III. INBOUND TRAFFIC ENGINEERING

One of the most challenging tasks today is shifting traffic between incoming link. This is because of the unpredictable nature of available controls: a downstream ISP cannot predict how much traffic will move without knowing the policies of the upstream. The task is even harder for transit ISPs than for edge ISPs because these controls can affect the volume of incoming traffic itself.

A BGP router in a transit AS generally receives several routes for a given prefix from its neighboring BGP routers, and each route is attached with various path attributes, such as LOCAL _PREF (local preference), AS PATH, and others, including proprietary attributes. The BGP router determines which route to accept based on the AS's import routing policy and attribute values. The route selection can be based on a highest LOCAL-PREF value, a shortest AS path length, e-BGP routes over i-BGP routes, and so on [6]. The AS path length is equal to the count of AS numbers in the AS path attribute. After determining the best route to a prefix, the BGP

router may further announce this route to a selected set of neighboring BGP routers but withhold it from another set, depending on the AS's export routing policy. As a result, different BGP routers end up having different views of the routes in the Internet. Moreover, without additional mechanisms, an AS cannot control the end-to-end forwarding path from an external source to a prefix inside the AS.

Incoming traffic can be influenced using modifying the MED attributes, selective announcements, prefix splitting, AS-path prepending, and BGP communities. Each of these approaches are explained below:

A. Modifying the MED attribute:

The first method to allow an AS to control its incoming traffic is to rely on the MED attribute. The MULTI_EXIT_DISC (MED) parameter is one of the attributes included in BGP protocol. MED is a 32-Bit integer value representing the non-transitive BGP metric. MED belongs to the set of BGP routing information parameters like NEXT_HOP, AS_PATH and LOCAL_PREF. MED parameter can be used to differentiate exit and entry points between two ASs. The MED parameter values are not advertised to other Autonomous Systems with the other route information [7]. This leads to the fact that MED is applicable balancing actor only in the situation of two or more connections between two neighboring Autonomous Systems (ASs). Typically this happens with two or more connections with the same ISP.

An ISP internal congestion may be exacerbated by its neighbors, because its neighbors might not be aware of the ISP's traffic engineering goals, internal topology, or load on internal links due to privacy reasons. Moreover, an ISP might not be willing to place such a high degree of trust in its neighbors. Hence, some mechanism to allow an ISP to control how much traffic it receives from each of its peering links is essential. Unfortunately, this is a highly challenging problem, as it requires the local ISP to influence route selection in remote ISPs, which in turn might wish to limit or completely ignore local ISP's goals. However, an ISP may convince its neighbour (through economic incentives) to allow the ISP to control how much traffic it receives on each link from the neighbor. This can be done by modifying the MED attribute, which can be used between a pair of ISPs connected via multiple peering links.

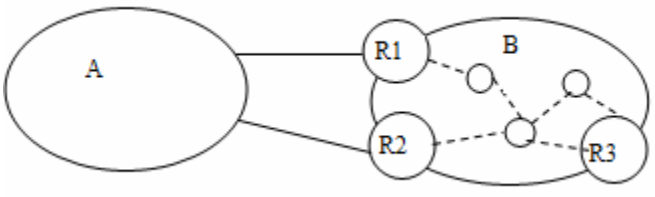

Figure 3. Example topology with two ISPs A and B.

For example, in figure 3, if B wanted to reduce the amount of traffic traversing through router R1, it could increase the value of MED attribute R1 advertises to A, causing the link to R2 to become more preferred by A's router and thereby decreasing R1's load.

Usage of MED in balancing inbound traffic is quite easy. However, it should be noted that the utilization of the MED attribute is usually subject to a negotiation between the two peering ASs, and some ASs do not take the MED attribute into account in their decision process. In the normal case MED values are not used when comparing route information received from different AS. With use of BGP's *always-compare-med* command this restriction is overruled and MED values are taken into account in every case. This can though lead to loops and route oscillation [7], [8].

B. Selective Announcement:

The second method that can be used to control the traffic that enters an AS is to rely on selective advertisements and announce different route advertisements on different links. This approach announces nonoverlapping prefixes to different links. For example, instead of announcing 158.32.0.0/16 to both upstream ISPs in Figure 4a, this approach announces two nonoverlapping longer prefixes to two different links. As a result, the traffic destined to these two prefixes will reach the network via the two respective links. Although this approach is very easy to deploy, it reduces the network resilience as only a single ISP is used for each prefix. Moreover, the actual AS path could be lengthened.

However, a drawback of this solution is that if the advertised link fails, the prefixes that were announced only on the failed link would not be reachable anymore.

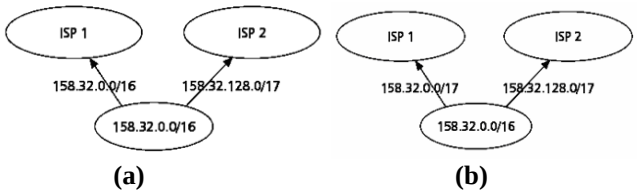

Figure 4. Static methods: a) Selective announcement, b) Prefix splitting

C. Prefix splitting:

A variant of the selective advertisements is the advertisement of more specific prefixes. Similar to the earlier approach, this method splits a prefix into longer prefixes. The difference is that the original prefix is also announced here. As shown in Figure 4b, the two prefixes advertised are overlapped, and the more specific one is sent to ISP 2. Under the longest-prefix-matching packet forwarding algorithm, traffic destined to 158.32.128.0/17 is expected to reach the network via ISP 2 only. Therefore, ISP 1 essentially serves as a backup for 158.32.128.0/17.

Clearly, this approach also suffers from the same problem of incurring a longer AS path. More important, the longer prefixes introduced by both approaches will cause BGP routing tables to grow very quickly. Because of that, BGP routers today are usually configured not to accept routes that exceed a certain prefix length (24 currently).

D. AS PATH Prepending

A BGP router's process to select the best routes from all accepted routes is complicated. A BGP router picks the route with the shorter AS Path among two equivalent routes after the comparison of their "local preference". Thus a possible way to influence the selection of the best routes by a distant AS is to *artificially* increase the length of the AS path by including multiple of its own AS number. This method, which is called *AS Path Prepending* (ASPP), is a popular BGP-based inbound traffic engineering method. In other words, a prepended AS path is an AS path that has some duplicated AS numbers that appear consecutively.

AS PATH is a mandatory attribute specified in BGP protocol specification. It contains a sequence of segment triples <path segment type, path segment length, path segment value> defining either ordered or unordered set of ASs that the UPDATE message has traversed [9].

Through ASPP, an AS could affect the distribution of traffic flowing into it. The usage of ASPP for inbound traffic engineering can be illustrated by the following example. Consider the traffic from AS_1 to AS_5 in Figure 2. In this network, AS_1 receives two routes for prefixes in AS_5: (AS_4,AS5) and (AS_5). These two routes have the same local preference because both of them are announced by AS_1's customer neighbors (AS_5 and AS_4), then the router in AS_1 selects the second route as the preferred route for prefixes in AS_5 since it has a shorter AS path. If AS_5 wishes that traffic from AS_1 goes through the link $AS_4 - AS_5$, it can use ASPP and announce AS path (AS_5,AS_5,AS_5) to AS_1. Now AS_1 receives two routes with AS path (AS_4,AS_5) and (AS_5,AS_5,AS_5). Therefore, the router in AS1 would choose the first route and its decision is changed.

This technique suits ideally in situation of end-user network but not so well for ISPs. This is because this new path length information is modified to external BGP peers and there to other ASs. Therefore in worst case this modification would lead to situation where no traffic flows through that ISP [8].

The effect of path prepending is explained with two different examples [11]: Consider the case where the multihomed network is connected to two ISPs that are similar: they interconnect at mostly the same Network Access Providers (NAPs) and Internet Exchanges, and they peer with mostly the same networks. Under these circumstances, other networks see two similar paths for the routes to announce. Figure 5 shows an example of this.

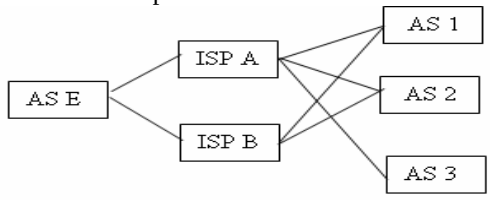

Figure 5. Multihoming to similar ISPs

If AS E is connected to two similar ISPs. But AS E wants to change the traffic flow and prepends the other path. In this situation there are three choices for traffic flow. Without path prepending traffic may or may not be balanced. When prepending path to ISP A, majority of the traffic flows through ISP B and when prepending path to ISP B, most of the traffic comes in over ISP A.

Table 1 shows which route is preferred in the situation shown in Figure 6 without path prepending, with prepending the path to ISP A, and with prepending the path to ISP B.

TABLE 1: PREPENDED PATHS OVER SIMILAR ISPS

	AS 1	**AS 2**	**AS 3**	**Traffic distribution**
Prepend to A	AEE **BE**	AEE **BE**	AEE **BE**	ISP A: 15% ISP B: 85%
No prepending	AE **BE**	**AE** BE	AE **BE**	ISP A: 40% ISP B: 60%
Prepend to B	**AE** BEE	**AE** BEE	**AE** BEE	ISP A: 90% ISP B: 10%

For the purposes of calculating the traffic distribution, it's assumed that A always handles 15% of the traffic, B always 10%, and AS1, AS2, and AS3 are all the source of 25% of incoming traffic. ASes with "odd" number (1, 3) prefer to send traffic over ISP B when the paths are of equal length; "even" ASes (2) prefer ISP A in this example. The preferred path is listed in bold type in the table.

When the two ISPs are not as similar, increasing the length of the AS path has a more gradual effect, because the paths over ISPs A and B aren't the same for all networks. Figure 6 shows multihoming to dissimilar ISPs.

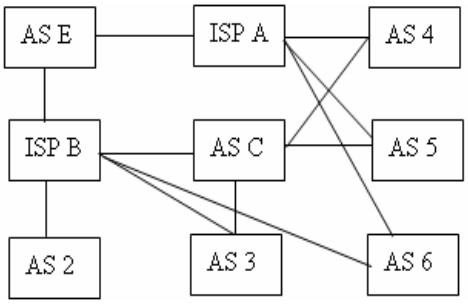

Figure 6. Multihoming to dissimilar ISPs

In this example, ISP B is a much smaller ISP that doesn't peer with networks AS4 and AS5, but rather buys transit service from AS-C to reach those networks. AS-C and AS3 do not have direct connection with ISP A they get service through ISP B. Network AS6, will immediately route traffic over ISP A when the path over ISP B is prepended, because the connections to both A and B are peering links. Table 2 shows the possible traffic distribution using prepending.

The traffic distribution in this example is 15% from ISP A, 5% from ISP B and ASes 2 and 3, 10% from AS C, and 20% from ASes 4, 5, and 6. It is a good idea to select dissimilar ISPs, for instance, one tier-1 ISP that peers with all the other large networks, and one tier-2 ISP that peers with many small networks. This way, you have a wide range of traffic engineering options.

TABLE 2: PREPENDED PATHS OVER DISSIMILAR ISPS

	ASC	**AS2**	**AS3**	**AS4**	**AS5**	**AS6**	**Traffic distrib--ution**
2 to ISP A	**BE**	**BE**	**BE** CBE	AEEE **CBE**	AEEE **CBE**	AEEE **BE**	A: 15% B: 85%
1 to ISP A	**BE**	**BE**	**BE** CBE	AEE **CBE**	**AEE** CBE	AEE **BE**	A: 35% B: 65%
No prepen ding	**BE**	**BE**	**BE** CBE	**AE** CBE	**AE** CBE	AE **BE**	A: 55% B: 45%
1 to ISP B	**BEE**	**BEE**	**BEE** CBEE	**AE** CBEE	**AE** CBEE	**AE** BEE	A: 75% B: 25%

E. BGP Communities

In addition to these above approaches, several ISPs have been using the communities attribute to give their customers finer control over the redistribution of their routes. The communities attribute is an optional variable size transitive attribute that can be attached to routes. This attribute can contain several 32 bits wide community values. First two octets have community attribute values coded with AS numbers and the other two octets are freely modified by the AS. A community is a group of destinations sharing a common property. BGP communities are special values that are attached to BGP advertisements.

Community values are often used to attach optional information to routes such as a code representing the city where the route was received or a code indicating whether the route was received from a peer or a customer. The community values can also be used for traffic engineering purposes. In this case, predefined community values can be attached to routes in order to request actions such as not announcing the route to a specified set of peers, prepending the AS-path when announcing the route to a specified set of peers or setting the local-pref**.** However, this technique relies on an ad hoc definition of communitv values and manual configurations of BGP filters, which makes it difficult to use and subject to errors.

The following traffic engineering actions are often supported:

1. do not announce the route : in this case the route with the associated community should not be announced to the specified peers.
2. prepend n times when announcing the route : the AS-path of the route with the associated community will be prepended n times when it is announced to the specified peers.
3. specify the value of the local preference to be used by the router that receives the route [10].

These actions typically apply toward a large AS (e.g. tier-1 or tier-2 ISPs providing transit service), an interconnection point, a country or a continent. Proper caution should be used when using communities since misconfigured community attribute might lead to a situation where no routes are advertised even if some of them would be wanted [11].

The Internet Engineering Task Force (IETF) is currentlv considering the definition of a new standard type of extended communities called *redistribution communities* [12] to solve the drawbacks of the utilization of classical communities to do traffic engineering. These redistribution communities can be attached to routes to influence the redistribution of those routes by the upstream AS. The redistribution communities attached to a route contain both the traffic engineering action to be performed and the BGP peers affected by this action. One of the supported actions allows an AS to indicate to its upstream peer that it should not announce the attached route to some of its BGP peers. Another type of action allows an AS to request its upstream to perform AS-PATH prepending when redistributing a route to a specified peer.

To understand the usefulness of such redistribution communities, let us consider again top part of the Figure 1, and assume that AS9 receives a lot of traffic from AS5 and AS6 and that it would like to receive the packets from AS5 (respectively AS6) via AS7 (respectively AS8) link. AS9 cannot ensure such a traffic distribution by performing AS-PATH prepending itself. However, this becomes possible with the redistribution communities by requesting AS7 to perform the prepending when announcing the AS9 routes to external peers. AS9 could thus advertise to AS7 its routes with a redistribution community that indicates that this route should be prepended two times when announced to AS6**.** With this redistribution community, AS7 would advertise path AS7:AS7:AS9 to AS6 and AS PATH AS7:AS9 to AS5. AS6 would thus receive two routes toward AS9, AS7:AS7:AS9 and AS8:AS9**,** and would select the route via AS8. AS5, on the other hand, would select the AS7:AS9 route that is shorter than the AS6:AS8:AS9 route.

IV. LIMITATIONS AND SUMMARY

The section above has described several techniques that can be used by ISPs to engineer their inbound traffic engineering. However, there are some limitations to be considered when deploying those techniques.

The MED attribute should only be used when there are multiple physical links between two ASes and not in the case of stub ASes multi-homed to several providers, a very common situation today.

An AS that announces the prefixes selectively on peering sessions does not guarantee connectivity to the prefixes when a session fails.

The control of the incoming traffic is based on a careful tuning of the advertisements sent by an AS. This tuning can cause several problems. An AS that advertises more specific prefixes or that divides its address space into distinct prefixes to announce them selectively will advertise a number prefixes longer than required. These prefixes will be propagated throughout the global Internet and will increase the size of the routing tables of all ASes in the Internet. [13] reports that more specific routes constituted more than half of the entries in a BGP routing table. To avoid this situation, several large ISPs have started to install filters to ignore the BGP advertisements corresponding to more specific prefixes [14].

The ASPP approach is often performed in a trial-and-error basis, since it is difficult to predict the outcome of performing AS PATH prepending on a given interdomain link in practice. The distributed prepending actions by different ASes may cause routing instability [15].

The redistribution communities can provide a finer granularity than ASPP and selective announcement. In practice, it can be expected that those communities will be used to influence the redistribution of routes toward transit ISPs having a large number of customers [16].

TABLE 3: SUMMARY OF APPROACHES USED FOR BALANCING INBOUND TRAFFIC.

Techniques	Scope	Conditions
MED attribute	Neighbors AS	Requires bilateral agreement
Selective announcement	Internet	Always works
Specific prefixes	Internet	If no filtering of long prefixes
AS-PATH prepending	Internet	Unsure
Community attributes	Internet	Always works

Table 3 provides a summary of various approaches used in balancing inbound traffic for multihomed ASs. Table gives the scope (local or remote) of each method and the condition under which the techniques work.

V. CONCLUSION

In this paper five approaches to balance inbound traffic are discussed. MED is easy to implement but had restricted functionalities. Announcing the selective advertisements on peering sessions on different links does not provide connectivity to the prefixes announced on the failed link (if a link fails). The longer prefixes introduced by both, selective announcement and prefix splitting approaches will cause BGP routing tables to grow very quickly. AS_PATH prepending is efficient but harder to implement since, it is difficult to select the appropriate value of prepending to achieve a given goal and needed better knowing of the network topology. BGP communities were sort of combination of everything. They are used to handle many functionalities like path prepending and how redistribution communities could allow an AS to flexibly influence the redistribution of its routes toward indirectly connected ISPs. All of the introduced tactics have their pros and cons. MED is suitable for small networks because of it's restrictions and easiness. Selective announcement and prefix splitting are easy to implement. Many large providers have implemented filters that reject advertisement for too long prefixes and avoid the growth of routing tables. AS PATH prepending is most suitable for situations where MED isn't enough. BGP Communities are the best choice when number of connections is large and many connections would have the same parameters. Then setting up communities to handle route advertising and traffic balancing is efficient.

VI. REFERENCES

[1] J. Hawkinson and T. Bates, "Guidelines for Creation, Selection, and Registration of an Autonomous System," RFC 1930, 1996.

[2] L. Gao, "On inferring autonomous system relationships in the internet," in *Proc. IEEE Global Internet Symposium*, Nov. 2000.

[3] B. Quoitin, S. Uhlig, C. Pelsser, L. Swinnen, and O. Bonaventure, "Interdomain traffic engineering with bgp," in *IEEE Communications Magazine*, May 2003, pp. 122-128.

[4] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771, Mar. 1995.

[5] F. Wang and L. Gao, "On inferring and characterizing Internet routing policies," in *Proc. ACM SIGCOMM Internet Measurement Workshop*, Oct. 2003.

[6] Jaggard, A.D., Ramachandran, V., "Robust Path-Vector Routing Despite Inconsistent Route Preferences", Proceedings of IEEE International Conference on network protocols, 2006, Nov. 2006, pp. 270 - 279

[7] T. Griffin and G. Wilfong, "Analysis of the MED oscillation problem in BGP." Paris, France: 10th IEEE International Conference on Network Protocols (ICNP'02), November 2002.

[8] Iljitsch van Beijnum., "Traffic Engineering: Local Routing Policy". O', www.onlamp.com/lpt/a/2796, 2002.

[9] P. Traina, D. McPherson and J. Scudder. Autonomous System Confederations for BGP. RFC 3065, IETF Network Working Group, February 2001.

[10] E.Chen and T. Bates. An Application of the BGP Community Attribute in Multi-home Routing. RFC 1998, IETF Network Working Group, August 1996.

[11] Iljitsch van Beijnum, "Traffic Engineering: Incoming Traffic" O'Reilly, OnLamp.com, www.onlamp.com/lpt/a/2797, 2002

[12] O. Bonahure, Stefaan D. C., J. Hass, and R. White, "Controlling the Redistribution of BGP Router," Internet draft. draft-ietf-ptomaine-bgp-redistribution-0l.txt. work in progress. Aug. 2002.

[13] A. Broido, E. Nemeth, and K. Claffy, "Internet expansion, refinement and ahurn", Eropean Transaction on Telecommunications, Jan 2002.

[14] S. Bellovin, R. Bush, T, Griffin and J. Renford, "Slowing routing table growth by filterning based on address allocation policies", available on www.research.att.com/ ~jrex, June 2002

[15] J. H. Wang, D. M. Chiu, John C. S. Lui and K. C. Chang, "Inter-AS Inbound Traffic Engineering via ASPP", IEEE Transactions on Network and Service Management, June 2007, pp-62-70.

[16] B. Quoitin, S. Uhlig, C. Pelsser, L. Swinnen, and O. Bonaventure, "A performance evaluation of BGP-based traffic engineering", Dec. 2004

VII. BIOGRAPHIES

Lata L. Ragha is now a Ph. D. candidate in the department of Computer Science and Engineering at the Jadapur University, Kolkata, India. She received her Bachelor and M.Tech degree in Computer Science and Engineering from Karnatak University in 1987 and Vishveshwarai Technological University, Karnataka, in 2000. She is currently working as Assistant Professor at Ramrao Adik Institute of Technology (RAIT), Mumbai University, Mumbai INDIA. Her research interests include Networking, Internet Routing, Routing Protocols, and Traffic Engineering.