Hot Potatoes Heat Up BGP Routing

Renata Teixeira
Computer Science & Engineering
U. California, San Diego
teixeira@cs.ucsd.edu

Aman Shaikh
Computer Engineering
U. California, Santa Cruz
aman@soe.ucsc.edu

Tim Griffin Jennifer Rexford
Internet & Network Systems
AT&T Labs–Research
{griffin,jrex}@research.att.com

Abstract— **The separation of *intra*domain and *inter*domain routing is a key feature of the Internet routing architecture. However, intradomain routing protocols such as OSPF and IS-IS *do* have a (sometimes significant) influence on the path-selection process in Border Gateway Protocol (BGP). In this paper, we argue that researchers should revisit the "interface" between the two tiers of the Internet routing system. Toward this end, we present an initial analysis of the impact of OSPF on BGP in a large ISP network. We propose a general methodology for associating BGP update messages with events visible in OSPF. Then, we apply our methodology to streams of OSPF link-state advertisements and BGP update messages. Our analysis shows that (i) "hot potato" routing is sometimes a significant source of BGP updates, (ii) BGP updates can lag 60 seconds behind the related OSPF event, which can cause delays in forwarding-plane convergence, (iii) OSPF-triggered BGP updates have a nearly uniform distribution across destination prefixes, and (iv) the fraction of BGP messages triggered by OSPF varies significantly across time and router locations, with important implications on external monitoring of BGP. Our measurement methodology and analysis results represent an important step in understanding the interplay between intradomain and interdomain routing.**

I. Introduction

The delivery of Internet traffic depends on the operation of the routing protocols running in and between thousands of Autonomous Systems (ASes). From the early days of the ARPANET, distributed management of the Internet by different institutions was an important design goal [1]. The Internet's two-tiered routing system allows ASes to exchange routing information without divulging their internal details to each other. This also allows each AS to select its own Interior Gateway Protocol (IGP). These IGPs are typically "metric based" to give operators control over resource allocation within their networks. For example, OSPF [2] and IS-IS [3] compute shortest paths based on link weights assigned by the operators. In contrast, the design of the Border Gateway Protocol (BGP) emphasizes scalability to a large number of ASes and address blocks, as well as the use of locally-configurable routing policies.

The separation of intradomain and interdomain routing offers other important performance benefits:

- **Scalability:** Link-state protocols such as OSPF and IS-IS rely on flooding of link-state advertisements (LSAs), which wouldn't scale to a network the size of the Internet.
- **Isolation:** The boundary between IGP and BGP decreases the influence of intradomain routing changes on the stability of the global Internet routing system.
- **Simplicity:** The two-tiered routing system would ideally enable a separation of concerns in reasoning about the properties of the IGPs and BGP, for debugging routing problems and analyzing measurement data.

These properties depend on the exact details of the interface between the two tiers of the routing architecture. We believe that the networking community needs to revisit these (sometimes implicit) design decisions, due to important changes in the past fifteen years:

- **Commercial constraints:** BGP routing policies depend on the commercial relationships between ASes, and "hot potato" routing causes a router to direct traffic to the "closest" exit point en route to the destination.
- **Network size:** Large ASes typically require a hierarchical distribution of reachability information via internal BGP (iBGP), resulting in more complex IGP/BGP interaction [4]. Also, large networks experience more topology changes due to failures and planned maintenance.
- **Traffic engineering:** Network operators change IGP link weights to adapt to changes in traffic and prepare for maintenance [5]. This may lead to more frequent BGP routing changes induced by "hot potato" routing.
- **Real-time applications:** Slow routing protocol convergence affects the performance of real-time applications such as telephony and gaming. Yet, ironically, the demands of these applications may encourage more frequent changes to the IGP weights to circumvent congestion.

In reality, the desirable properties of scalability, isolation, and simplicity do *not* necessarily hold, since the IGPs affect BGP in terms of:

- **Hot-potato routing:** The IGP path distances affect the BGP decision process. If multiple BGP routes are "equally good," the router selects the route with the "closest" exit point in the IGP sense.
- **Next-hop reachability:** The IGP determines whether the routers in the AS believe that the exit point associated with the BGP route is reachable.
- **iBGP message delivery:** The iBGP sessions used to propagate BGP routes inside the AS depend on the IGP for message delivery. Transient packet loss during IGP routing convergence could trigger iBGP session failures.
- **Multi-exit discriminator:** An AS can use the MED attribute in BGP to specify exactly where traffic should enter the AS. When used, MED metrics are often tied directly to the AS's IGP path distances, making internal IGP instability visible to neighboring ASes[1].

Our primary goal is to understand the interplay between IGPs and BGP through measurement studies of operational networks. After a brief background discussion of IP routing in Section II, Section III proposes a methodology for combining independent measurements of the OSPF and BGP protocols to identify which BGP update messages are triggered by OSPF events. Then, Section IV presents our initial measurement results from a joint analysis of OSPF LSAs and BGP update messages collected from the IP backbone of a large ISP. Although previous mea-

[1] We do not investigate MED-based *cold-potato* routing in this paper.

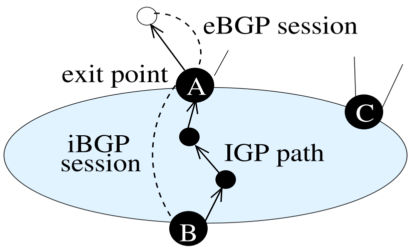

Fig. 1. Interaction between eBGP, iBGP, and IGP

0. Ignore if exit point unreachable
1. Highest local preference
2. Lowest AS path length
3. Lowest origin type
4. Lowest MED (with same next-hop AS)
5. eBGP-learned over iBGP-learned
6. Lowest IGP path cost to exit point ("Hot potato")
7. Lowest router-id of BGP speaker

TABLE I
STEPS IN THE BGP DECISION PROCESS

surement studies have characterized OSPF/IS-IS LSAs [6–9] or BGP updates [9–12] in isolation, we believe this is the first paper to present a joint analysis of the two datasets. The paper concludes in Section V with a summary of our results and a discussion of avenues for future work.

II. Operational View of IP Routing

In large service provider networks, the forwarding table at each router depends on the interaction between multiple routing protocols. In the example in Figure 1, router A learns a route to a destination prefix from an external BGP (eBGP) neighbor and propagates this information to router B via internal BGP (iBGP); B uses IGP information to (i) determine that A is the closest exit point and (ii) compute the outgoing link along a shortest path to A. This section explains these protocols and their interactions in more detail, focusing on OSPF as an example of an IGP.

A. Open Shortest Path First (OSPF)

OSPF [2] is a link-state routing protocol where each unidirectional link is assigned an administrative weight. The reliable flooding of link-state advertisements (LSAs) ensures that each router can construct a complete view of the network topology. LSAs are flooded periodically and in response to network events, such as weight changes and equipment going up or down. Each router runs Dijkstra's algorithm to compute the shortest paths to every other node and uses the results to build the forwarding table. This ensures that, in steady state, each IP packet is forwarded along a shortest path in terms of link weights. For scalability, OSPF allows the network to be divided into areas to define a two-level hierarchy. Area 0 (the backbone area) resides at the top level of the hierarchy and provides connectivity to the other areas.

B. Border Gateway Protocol (BGP)

BGP [13] is a path-vector protocol that allows each AS to apply local policies in selecting and propagating routes for each destination prefix. Two routers exchange BGP messages over an underlying TCP connection. BGP routers send new update messages only when something has changed. An *advertisement* notifies a neighbor of a new or a modified route, whereas a *withdrawal* revokes a route that is no longer available. An advertisement may be a "replacement" of an earlier route (i.e., an *implicit withdrawal*) or a new "announcement" for a prefix. Each advertisement includes various route attributes, including the list of ASes along the path.

A large backbone network typically has multiple BGP-speaking routers, and BGP sessions with multiple neighboring ASes. Such a network can also have multiple BGP sessions with each neighbor AS. As a result, a router may receive routes for a destination prefix from multiple neighbors. The router applies *import policies* to filter unwanted routes and to manipulate the attributes of the remaining routes. The router then invokes a *decision process* to select exactly one "best" route for each destination prefix among all the routes learned from its neighbors. Different routers in an AS apply the BGP decision process independently and might select different "best" routes for the given prefix, depending on their locations in the network. Ultimately, each router applies *export policies* to manipulate attributes and decide whether to advertise the best route to each neighbor.

In addition to having BGP sessions with neighboring ASes, routers may use BGP to distribute routing information within an AS. Instead of having a full mesh of iBGP sessions, a large AS may introduce hierarchy through the use of route reflectors or confederations [14].

C. OSPF Impact on BGP Updates

Table I summarizes the steps in the BGP decision process. Several steps depend on BGP attributes (such as local preference, AS path length, origin type, and MED) that are conveyed in route advertisements and can be manipulated by local policies. However, OSPF controls the two steps listed in bold-face. First, the router must determine if the BGP "next hop" (the "exit point") is reachable. Then, if multiple routes proceed through the next five steps of the decision process, the OSPF path distance is used to select a route with the nearest exit point. For example, in Figure 1, the failure of router A, a link failure inside the network, or a change in OSPF weights could cause router B to select the route from C.

To summarize the interaction between OSPF and BGP, consider what happens when a link fails along the path from B to A, causing C to become the closest exit point:
1. **LSA flooding:** One (or both) of the end-points of the link detect that a failure has occurred and initiates reliable flooding of an LSA to the rest of the network.
2. **OSPF processing:** B receives the LSA, updates its link-state database, and recomputes its shortest paths.
3. **BGP processing:** B revisits the BGP routing decisions for each destination prefix. Some decisions come down to step 6 that depends on the OSPF path distance. Router B selects a new best path with exit point C for these prefixes and modifies its BGP routing table. Then, B combines the OSPF and BGP information to modify the forwarding-table entries for each of these prefixes.
4. **BGP updates:** B considers whether to send an update message to its BGP neighbors. B would send the new route to iBGP

neighbors, such as route-reflector clients if B is a route reflector. For each eBGP neighbor, B applies the export policy to filter the route or modify its attributes. B exports the route if any externally-visible attributes have changed (e.g., the new route has a different AS path, perhaps with the same length).

Even if eBGP neighbors do not receive a new update message, changes in the exit point affect how traffic flows through the network and on to neighboring ASes.

III. Analysis Methodology

In this section, we first describe how we process a stream of OSPF LSAs to identify "events" that may affect BGP. Next we discuss how we classify the BGP update messages to associate them with related OSPF events.

A. BGP-Visible OSPF Routing Events

OSPF affects BGP through changes in the reachability or path distance from one router to another. Some OSPF link-state advertisements (LSAs) have no influence on BGP: (i) LSAs sent periodically as part of OSPF soft-state refresh, (ii) links with high OSPF weights that do not appear on any shortest path (e.g., links under maintenance or provisioning), and (iii) links that always appear as part of multiple shortest paths with other links (e.g., parallel links between two routers). Other OSPF LSAs may affect the BGP decisions at *multiple* routers in the network. For example, a single link failure or OSPF weight change might alter the shortest-path distance for multiple pairs of routers. A single router coming up or down would change the exit-point reachability for routers throughout the AS.

Given a stream of OSPF LSAs, we identify individual "events" that can affect BGP:

1. **Change LSAs:** The algorithm filters OSPF LSAs originated as part of periodic refresh to focus only on LSAs triggered by network changes.
2. **Path computation:** For each "change LSA", we emulate the OSPF shortest-path computation [2] to determine the distance from each router to every other router.
3. **BGP-visible changes:** For each router, we identify any changes in the path distance for reaching other routers. The output of this step is zero or more entries that were added (ADD), deleted (DEL), or changed (CHG). Each entry includes a time stamp and the loopback IP addresses of the source and destination routers; ADD and CHG entries also include the path distance of the (new) shortest path to the destination router.

In the next subsection, we identify cases where the ADD, DEL, and CHG events may be responsible for a change in the source router's BGP routing decisions.

B. BGP Update Classification

Starting with an initial BGP routing table, a stream of BGP update messages from a router reflects a series of changes in the best route for certain prefixes. Determining why a router changes from one route to another is difficult in practice, since multiple events may trigger the same BGP update message. The root cause is not necessarily an OSPF event. New advertisements from neighboring ASes and changes in local routing policies can cause changes in the best path as well. For example, suppose that a router switches its best route for a destination prefix from r to s. One possibility is that an eBGP neighbor started advertising a more attractive route s (say, with a shorter AS path); another possibility is that an increase in the IGP path distance to r's exit point caused the router to prefer a closer exit point with route s.

To aid the analysis, we propose a classification of BGP update messages that identifies the types of OSPF events that *could* explain a change we see in the BGP-level routing decision. Figure 2 illustrates how we classify a BGP update message for a prefix (p) generated by a particular router (X). For each class of BGP update, the figure also lists the types of possible OSPF events that could have caused it. Let us explain the classification of Figure 2 by following down the branches in the decision tree.

We refer to the best route for the prefix before the update as r, and the best route after the update as s. If r is null, the figure categorizes the update as an *announcement* for prefix p. The possible cause is an OSPF ADD that made the exit point $s.exit$ reachable from the router. Similarly, if s is null, the update is a *withdrawal* of the route r. The possible cause in this case is a DEL that made $r.exit$ unreachable.

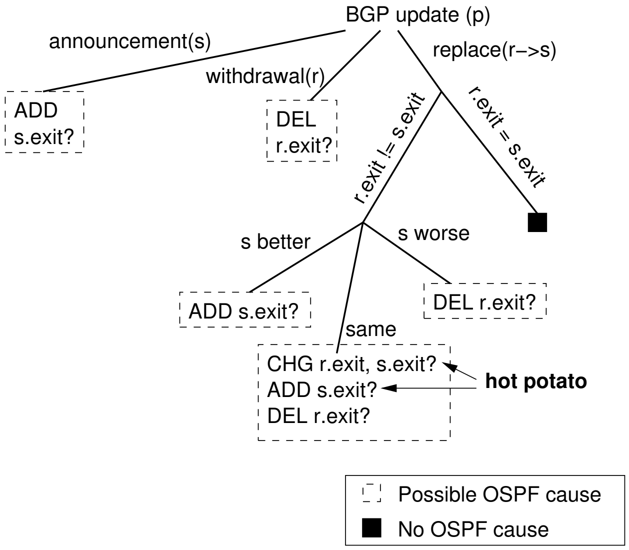

Fig. 2. Classification of BGP updates sent by a single router

If both r and s are non-null, route s *replaces* r. For this case, further analysis narrows down the possible OSPF causes. If the two routes have the same exit point, then OSPF cannot have triggered the routing change; indeed, this is the only case where we can completely dismiss OSPF as a possible cause simply by examining the BGP message. When the exit points differ, a more detailed comparison of r and s in terms of steps 0–5 in the BGP decision process in Table I refines our understanding of the possible role of OSPF. An ADD that makes $s.exit$ reachable could allow a switch to a better route, whereas a DEL that makes $r.exit$ unreachable could force a switch to a worse route. However, if the two routes are equally good, the change could be caused by either of these two events, or by a CHG that increases the distance to $r.exit$ or decreases the distance to $s.exit$.

C. Matching BGP Updates and OSPF Routing Changes

We associate BGP update messages with specific OSPF events by emulating the changes in the OSPF and BGP routing state from the viewpoint of a particular router. Starting with initial OSPF and BGP routing tables, our algorithm attempts to match each BGP update with a specific OSPF event. For each BGP update message U, we use the following criteria to determine whether an OSPF event E can be considered as a candidate:

1. The type of E should match the classification scheme of Figure 2.
2. The source and destination routers of E should be the same as source router and exit point of U, respectively.
3. E should occur "close" in time relative to the BGP update message U.

For example, a BGP announcement of route s at time 5 could be matched with an ADD of $s.exit$ that happened at time 4; however, the BGP announcement would not be matched with DEL or CHG events, no matter when they occurred or what exit points they involved. If a replacement of r by s appears to match a CHG event, we verify that the OSPF change makes $s.exit$ closer than $r.exit$, or makes $r.exit$ further than $s.exit$.

The algorithms faces two main challenges:

- **Timing issues:** In some cases, a BGP message does not match *any* related OSPF event. In other cases, multiple matches may occur at different points in time, and we select the nearest match within a specified time window. In fact, with a very large time window, our algorithm would probably find a relevant OSPF event for nearly every BGP update message! On the other hand, a very small time window would lead to false negatives, and an undercounting of the influence of OSPF on BGP. In practice, *negative* delays between the OSPF event and the BGP update message could conceivably happen, depending on where the data are collected. The reliable flooding of OSPF LSAs is typically implemented in software on the router, which may subject these messages to higher delays. In contrast, BGP updates are sent via a TCP connection between two routers; the IP packets carrying these messages traverse the "fast forwarding path" through the network. As such, BGP messages might reach a measurement location before the OSPF LSA responsible for the routing change.
- **eBGP Effect:** Even with a "good" setting of the time window, our algorithm could match a BGP message with an *unrelated* OSPF event. The BGP routing change could be caused by an external event, such as a policy change or an eBGP route update (e.g., a withdrawal of a route r by an eBGP neighbor that causes the router to replace r by s). Yet, a seemingly-related OSPF event could occur nearby in time (e.g., a CHG event that makes $s.exit$ closer than $r.exit$). Our algorithm would mistakenly associate the replacement of r by s with the OSPF event. (In practice, the OSPF event would have caused a similar routing change anyway if the BGP event hadn't happened first!) In any case, these kinds of mismatches are difficult to avoid. Our fine-grained classification of the BGP routing changes in Figure 2 reduces the possibility of these kinds of mismatches. Careful selection of the time window helps as well, as we explore experimentally next.

IV. Measurements and Results

We applied our matching algorithm to OSPF LSAs and BGP update messages collected in a large commercial backbone. After a brief description of our measurement infrastructure, we show that BGP updates typically occur within one to two minutes of the related OSPF event, due perhaps to the 60-second BGP "scan" timer [15]. Then, we show that OSPF-triggered BGP updates have a nearly uniform distribution across destination prefixes, in sharp contrast to the remaining BGP updates. Last, we discuss how the impact of OSPF events on BGP routing changes varies significantly across time and router location.

A. *OSPF and BGP Monitoring*

We have deployed OSPF and BGP monitors in the commercial backbone to collect OSPF LSAs and BGP updates. The OSPF monitor establishes an adjacency with a router in area 0 of the backbone, and archives all LSAs. The BGP monitor has an iBGP session (running over TCP) to the same router; using an *iBGP* session allows the monitor to see changes in the "exit point" of BGP routes. The BGP monitor also dumps its own routing table once a day to provide an initial view of the best route for each prefix. To minimize the timing effects, both monitors run on the same server, which has a physical connection to the router. The analysis in this section focuses on data collected in June 2003. Neither monitor experienced any disruption in its routing session during this period.

To understand if IGP events cause iBGP session resets, we analyzed data collected by a separate monitor with iBGP sessions to routers throughout the network. These iBGP sessions include multi-hop TCP connections to remote locations that would be more vulnerable to disruption during IGP routing convergence. During June 2003, each of these iBGP sessions experienced 2–3 resets, perhaps due to temporary disruption of the monitor's connection to the rest of the network. These results suggest that OSPF events were not a significant contributor to iBGP session resets in the network. In analyzing the OSPF event stream in isolation, we find very few ADD and DEL events. This is not surprising since ADD and DEL events would only occur when routers go up or down (say, due to failures or router reboots) or the network becomes partitioned. Ultimately, nearly all OSPF events reflect changes in path distance.

B. *Time Difference Between OSPF and BGP*

We first explore the timing of OSPF-triggered BGP updates, with the goal of identifying a good window size for matching BGP updates with OSPF events. Our analysis focuses on data collected on June 25, 2003. We discuss the daily variation of these results in Section IV-D. In Figure 3, we plot the cumulative distribution of the time between the receipt of a BGP update message and the nearest OSPF LSA identified by our matching algorithm. For completeness, we consider OSPF events from ten minutes before the BGP update to ten minutes after. However, the x-axis in Figure 3 shows the period for BGP updates that arrive between ten seconds before the OSPF LSA to three minutes after, where the vast majority of the matches occur. The y-axis plots the cumulative percentage of *all* BGP updates associated with an OSPF event. Some BGP messages are not matched with any OSPF event. Our algorithm found matches for just 14.8% of the BGP updates, even with a 20-minute time window.

The graph shows that most of the OSPF-triggered BGP messages are generated within 60 seconds of the receipt of the re-

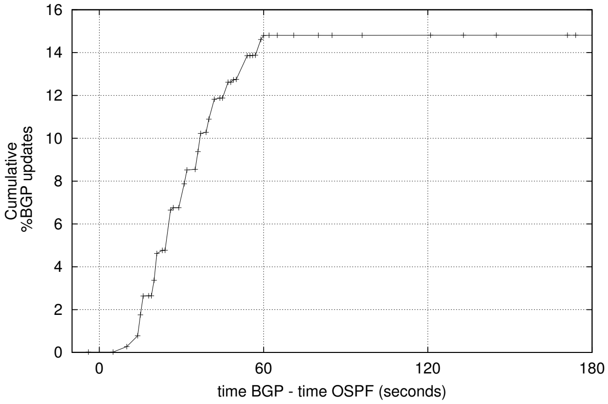

Fig. 3. Time difference between BGP update message and OSPF event

lated OSPF LSA. The shape of the graph can be explained in terms of several low-level timers controlling the operation of the routing protocols. Cisco routers have timers (spf-delay and spf-holdtime [16]) that typically introduce 5–10 seconds of delay in starting the shortest-path computation after an LSA arrives. In addition, BGP routes are not updated immediately after an OSPF routing change. Instead, a scanner process runs every 60 seconds to update BGP routes [15]. Since LSAs can arrive at any time within this 60-second period, we observe that the cumulative percentage of BGP updates increases linearly within the 0–60 second window in Figure 3.

The relatively long delays between the OSPF event and the BGP routing change may have important implications on the convergence of the forwarding plane[2]. Convergence delay has been a subject of much interest in the networking research community in the past few years. However, previous work has focused on OSPF/IS-IS [17] and BGP [18, 19] in isolation. The results in Figure 3 suggest that the interaction between the two protocols may have important implications on the convergence of the forwarding path.

Although the x-axis in Figure 3 covers times in the $(-10, 180)$ range, our algorithm did match a small number of BGP messages to OSPF events over larger time windows. We believe that these are false matches caused by two independent events (as explained in Section III-C). Based on the results in Figure 3, and similar results on different days, we believe that a time window of $(-10, 180)$ is a reasonable way to reduce the number of false matches; we use this window for the rest of this section. In our ongoing work, we are exploring additional ways to avoid incorrect matches.

C. Impact of OSPF on BGP Updates Across Prefixes

Previous work has shown that a small fraction of prefixes are responsible for most of the BGP route updates [9]; more recent studies have shown that the popular prefixes responsible for most of the traffic experience very few BGP updates [11, 12]. Figure 4 plots the cumulative distribution of BGP update messages versus the cumulative percentage of the prefixes for the June 2003 data. The prefixes are sorted according to their contribution to the number of BGP messages. The middle curve corresponds to all the BGP messages. About 20% of prefixes contribute 70% of the BGP updates, consistent with previous findings.

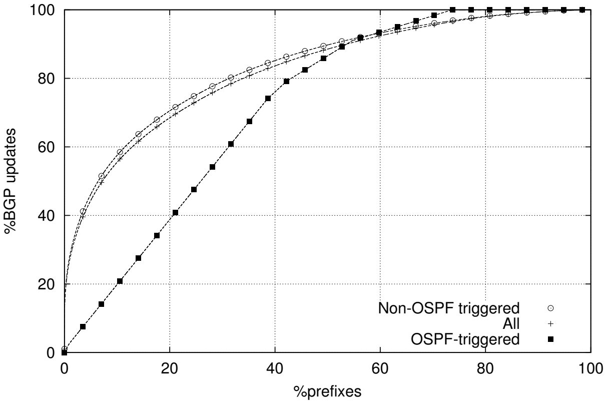

Fig. 4. Cumulative distribution of BGP updates across prefixes

However, the bottom curve shows that the distribution of *OSPF-triggered* BGP updates is nearly uniform across the destination prefixes! This occurs because OSPF CHG events can affect the path costs to reach exit points across a wide variety of prefixes. Still, some prefixes do not experience many OSPF-triggered BGP updates, as seen in the flat portion in the upper-right part of the graph. We believe this corresponds to prefixes with only one BGP exit point; "hot potato" routing is not relevant for prefixes with a single exit point. Still, the uniform distribution across the bulk of the prefixes may have important implications. For the most stable prefixes, OSPF events may be the primary cause of iBGP routing changes. Since some of these prefixes are responsible for a large amount of traffic, limiting the changes in how traffic travels to neighboring domains could be quite useful.

D. Variability over Time and Router Location

The influence of OSPF on BGP varies significantly across time. Figure 3 shows that OSPF events were responsible for about 14.8% of the BGP updates on a single day in June. However, this number varied from 0% to 23% over the 30 days in the month: this number was less than 1% for 23 of the 30 days. In general, the vast majority of the OSPF CHG events did not affect the proximity of the exit point for any BGP prefix at this router.

However, the likelihood that a CHG event affects the selection of the BGP best route depends on how close the router is to each of the exit points. For a router in the same Point-of-Presence (PoP) as one of the exit points, the probability that an OSPF CHG event would make another exit point more attractive is extremely low. To study this effect, we analyzed the data collected by the route monitor with separate iBGP sessions with multiple route reflectors in the network. We applied our matching algorithm to this data using the OSPF LSAs collected at the other monitoring location, with the clocks on the two monitors

[2] Note that the extra delay for BGP routing changes does *not* affect the stability of the forwarding path for the iBGP sessions. The IP packets sent over iBGP sessions travel between routers within the backbone, and the forwarding of traffic between these routers depends only on OSPF!

synchronized using NTP. The fraction of BGP updates triggered by OSPF events per day varied from 0–47% across the route reflectors for the month of June.

These results may have important implications on external BGP monitoring at public collection points like RouteViews [20]. Depending on which router in an AS provides the feed to RouteViews, the external data may look very different. The external view of the data would vary depending on what fraction of the internal OSPF events trigger BGP routing changes, and what fraction of these BGP changes are exported via eBGP.

V. Conclusions

The interplay between intradomain and interdomain routing has important implications on the stability and efficiency of Internet routing and, in turn, on end-to-end performance. In this paper, we have presented a methodology for joint analysis of OSPF and BGP measurement data and an initial characterization of the interplay between the protocols. Our initial results suggest that hot-potato routing may play an important role in BGP routing changes, and that BGP updates can lag 60 seconds behind the related OSPF events. We also show that the fraction of BGP updates that are triggered by OSPF events varies significantly across time and router location, suggesting a need for further analysis and modeling of how the protocols interact. Our ongoing work focuses on:

- **Matching algorithm:** To improve the matching of BGP updates with OSPF events, we plan to correlate BGP updates from multiple vantage points in the network and group updates appearing close together in time. We also want to extend our classification scheme to incorporate the details of how OSPF-triggered BGP updates are propagated through the iBGP hierarchy.
- **Data analysis:** We plan to explore how much influence the OSPF-triggered BGP updates have on the flow of traffic in the network, by combining our analysis with flow-level measurements of the traffic at the prefix level. We are also investigating what fraction of OSPF-triggered BGP updates are exported to eBGP neighbors and how much these kinds of updates might affect the analysis of BGP measurements collected at public route servers.
- **Improving isolation:** We are investigating protocol extensions and operational practices that would decrease BGP's reaction to IGP changes. For example, protocols such as MPLS could be used to tunnel to exit points, to separate "hot-potato" routing from the details of the path distances between routers. We are also investigating ways to tune IGP weights (for traffic engineering and planned maintenance) without triggering BGP routing changes.

References

[1] D. D. Clark, "The design philosophy of the DARPA Internet protocols," in *Proc. ACM SIGCOMM*, pp. 106–114, August 1988.

[2] J. Moy, "OSPF Version 2." RFC 2328, April 1998.

[3] R. Callon, "Use of OSI IS–IS for Routing in TCP/IP and Dual Environments." RFC1195, December 1990.

[4] G. Wilfong and T. G. Griffin, "On the correctness of IBGP configuration," in *Proc. ACM SIGCOMM*, August 2002.

[5] B. Fortz, J. Rexford, and M. Thorup, "Traffic engineering with traditional IP routing protocols," *IEEE Communication Magazine*, October 2002.

[6] D. Watson, C. Labovitz, and F. Jahanian, "Experiences with Monitoring OSPF on a Regional Service Provider Network," in *Proc. International Conference on Distributed Computing Systems*, pp. 204–213, May 2003.

[7] A. Shaikh, C. Isett, A. Greenberg, M. Roughan, and J. Gottlieb, "A Case Study of OSPF Behavior in a Large Enterprise Network," in *Proc. Internet Measurement Workshop*, November 2002.

[8] G. Iannaccone, C. Chuah, R. Mortier, S. Bhattacharyya, and C. Diot, "Analysis of link failures in an IP backbone," in *Proc. Internet Measurement Workshop*, November 2002.

[9] C. Labovitz, A. Ahuja, and F. Jahanian, "Experimental Study of Internet Stability and Wide-Area Network Failures," in *Proc. International Symposium on Fault-Tolerant Computing*, June 1999.

[10] C. Labovitz, R. Malan, and F. Jahanian, "Internet Routing Instability," *IEEE/ACM Trans. Networking*, vol. 6, pp. 515–558, October 1998.

[11] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "BGP Routing Stability of Popular Destinations," in *Proc. Internet Measurement Workshop*, November 2002.

[12] S. Agarwal, C.-N. Chuah, S. Bhattacharyya, and C. Diot, "A Study of the Impact of BGP Dynamics on Intra-Domain Traffic," Sprint ATL Research Report RR03-ATL-051677, Sprint ATL, May 2003.

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

[14] S. Halabi and D. McPherson, *Internet Routing Architectures*. Cisco Press, second ed., 2001.

[15] Understanding BGP Processes on Cisco. http://www.cisco.com/warp/public/459/highcpu-bgp.html#topic1.

[16] "Configure Router Calculation Timers." http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/%np1_c/1cprt1/1cospf.html#xtocid2712621.

[17] C. Alaettinoglu, V. Jacobson, and H. Yu, "Toward milli-second IGP convergence." Expired Internet Draft, draft-alaettinoglu-isis-convergence-00.txt, November 2000.

[18] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet routing convergence," *IEEE/ACM Trans. Networking*, vol. 9, pp. 293–306, June 2001.

[19] Z. M. Mao, R. Govindan, G. Varghese, and R. Katz, "Route Flap Damping Exacerbates Internet Routing Convergence," in *Proc. ACM SIGCOMM*, August 2002.

[20] "Route Views Project." http://www.routeviews.org.