

Link-Rank: A Graphical Tool for Capturing BGP Routing Dynamics

Mohit Lad, Lixia Zhang
Computer Science Department
University of California
Los Angeles, CA 90095, USA
{mohit,lixia}@cs.ucla.edu

Dan Massey
USC/Information Sciences Institute
3811 N Fairfax Dr.
Arlington, VA 22203, USA
masseyd@isi.edu

Abstract

Failures at the BGP level can have significant impact on the overall Internet. Understanding the behavior of BGP is thus both an important practical challenge and an interesting research problem. To understand the true dynamics, and help interpret the multiple gigabytes of BGP log data, we have developed the “Link-Rank” graphical toolset. Link-Rank weights the links between Autonomous Systems by the number of routing prefixes going through each link. Tracing these graphs over time results in a directed graph that shows the weight changes of the logical inter-AS links. From this graph one can easily visualize the complex BGP path changes and also combine views from multiple vantage points, to get a better picture of global routing dynamics. We illustrate the usefulness of Link-Rank by using it to examine BGP routing dynamics in three example cases. These examples show that Link-Rank is able to help BGP analysts estimate the scope of routing changes and to reveal important routing dynamics in the presence of superfluous BGP update messages.

Keywords

Border Gateway Protocol(BGP), BGP Routing Dynamics, Network and Systems Monitoring

1. Introduction

The Internet consists of a large number of Autonomous Systems (ASes) and the Border Gateway Protocol (BGP) [1] is used to exchange reachability information between Autonomous Systems. In order to monitor global reachability, analyze BGP dynamics, and promptly detect routing failures, many individual ASes monitor their local view of the BGP routing system; some ASes also provide public access to local routing table snapshots, referred to as *looking glasses*. A few passive BGP monitoring sites, such as RIPE [2] and RouteViews [3], have also been established to collect BGP update data. These monitoring sites peer with BGP routers of various ASes to passively collect the updates from the peering routers, and make the resulting BGP log data publicly available.

Although publicly accessible looking glasses and BGP update logs provide potentially useful information into the operations of the Internet routing infrastructure, we face several challenges in making effective use of the available data. First, due to routing policies, each vantage point tends to have a different view of BGP reachability and routing activities, and observations at a particular router do not reflect the state for the rest of the

Internet in general. Second, there is no easy or clear way to combine these individual views into a coherent picture of the global routing changes. Third, the data volume is large. For example, Route-Views typically logs over 10 gigabytes of BGP update messages per month from over 30 routers in various topological locations; during the SQL Slammer worm attack on Jan 25th 2003, RIPE monitoring point received 632795 updates from a single peer router. [4] also showed that BGP log data collected during stressful events, such as Nimda and Code Red worm attacks, may contain a high percentage of measurement artifacts. The sheer amount of log data makes it difficult, if not impossible, to observe whether the global routing system is going through its daily routine reachability exchanges, or a significant failure has occurred, and if so where it is, what is the scope of the impact, and how well BGP is adapting to the failure.

As the Internet continues to grow as measured by the number of ASes, the increase in inter-AS connectivity, the variety of routing policies, and the lack of monitoring tools makes it increasingly challenging to understand BGP routing dynamics. In order to provide dependable Internet connectivity, we not only need to collect BGP operation logs, but we must also have effective tools to analyze and interpret the log data to explain the observed routing update dynamics, and present the data in a way that makes the global routing dynamics readily accessible to human operators.

In this paper we present the *Link-Rank* toolset. We define Link-Rank graphs showing the weighted importance of AS peerings from observation points. We also define Rank-Change graphs that can extract a compact graphical representation from BGP log data to visually show important BGP dynamics. We further provide a systematic approach for combining the Rank-Change graphs from multiple routers into a single global graph summarizing overall BGP dynamics. The global graph can help in understanding issues like the nature of the event and extent of overlap of affected peerings. But unlike the efforts on constructing the AS level topology of the Internet [5, 6, 7], the Link-Rank objective is to present BGP *routing dynamics* in an aggregate and easy to understand form in order to aid network operators in understanding and evaluating the BGP operations. We use three case studies to show how Link-Rank can compactly represent bulk BGP log data. In the first case, RouteViews vantage point observed a burst of routing changes to around 10000 prefixes, from a single router. Using Link-Rank, we show that this event is local, and does not affect global routing. In the second example, we apply Link-Rank to BGP logs collected during the time period of a known failure at a large AS, and show how Link-Rank helps assess the global impact of this failure. Finally, we use Link-Rank to examine the global routing events during the SQL Slammer worm attack on January 25, 2003.

The paper is organized as follows. Section 2 formally presents the Link-Rank model and presents the algorithm for graph construction and combining the views from multiple monitoring points. Section 3 show how Link-Rank is used to analyze BGP routing events, and provides a detailed look at a routing update burst, and a large AS failure from an individual peer's point of view. Section 4 presents another application of Link-Rank to examine routing during the SQL Slammer worm attack. Finally, Section 6 concludes the paper.

2. Link-Rank and Rank-Change graphs

This section defines *Link-Rank* and *Rank-Change* graphs that concisely represent large-scale path vector dynamics in a simple form and explains them with the help of simple examples. In a path vector routing protocol such as BGP, routers exchange update messages that include the complete path used to reach the destination and the path information is typically used in loop detection, but in this case we use it for capturing BGP routing dynamics.

2.1 Definition and Construction of Link-Rank Graphs

A Link-Rank graph is constructed from a path vector routing table by representing the paths as a directed graph. Each link in the graph is assigned a weight (i.e. rank) in terms of the number of routes that rely on the link. Figure 1(a) shows a sample topology (shown by the undirected edges). In this example topology, there are three prefixes that are advertised: prefix A is advertised by AS5, prefix B is advertised by AS6, and prefix C is advertised by AS7. There is a BGP peering session between our observation point and a router in AS1 so the path information is sent to the observation point. To reach prefix A, the observation point's next hop is AS1, followed by AS2 and then finally AS5, that is the origin of the prefix.

The objective of a Link-Rank graph is to represent the AS level links used from a particular vantage point in a visually compact form. Later, we describe how we use Link-Rank graphs to construct Rank-Change graphs for capturing dynamics. The bold directed edges in figure 1(a) constitute the Link-Rank graph for point X, showing how the three prefixes can be reached from AS1. The weight on each link indicates the number of prefixes that are reached by using the link. Note that the Link-Rank graph from a point will include only those ASs and links that are used to reach prefixes. In this example, we have presented the underlying topology, but in actual analysis, one does not know the complete AS-AS peering or connectivity.

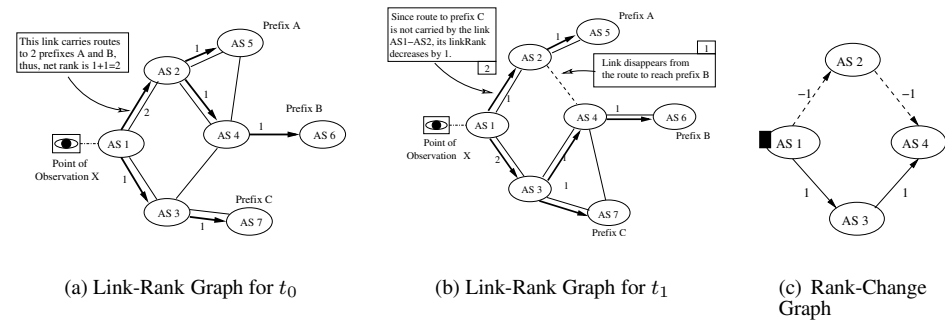


Figure 1 Graphs from Point X

More formally, we define the Link Rank graph as follows:

Given a path vector routing table with entries of the form $\langle prefix_i, path_i \rangle$, the corresponding *Link Rank* graph is a directed graph $G = (V, E)$ where :

- $V = \{v | v \in path_i \text{ for some } i\}$.
In other words the vertices in the Link Rank graph are the nodes from the paths.
- $E = \{\langle v, w \rangle \text{ such that } \langle v, w \rangle \text{ is a subsequence in some } path_i\}$.
In other words, the set of edges in the Link Rank graph correspond to the links in the paths.

Each link in the Link-Rank graph is assigned a weight based on the *importance* of the link. The measure of importance used in this study is number of routes that rely on that link, but other measures like address space coverage can also be used. More precisely, $rank(\langle v, w \rangle \in E) =$ the number of paths, $path_i$, that contain subsequence $\langle v, w \rangle$.

2.2 Rank-Change Graphs: Tracing Link-Rank Changes Over Time

A Link-Rank graph from a vantage point provides a static snapshot of the AS-AS peerings used to carry routes. Rank-Change graphs, which are derived over time from Link-Rank graphs, help in visually identifying the sections of the routes that change. Conceptually, a Rank-Change graph denotes the change in BGP routes as observed from a router over a period Δt . The Rank-Change graph shows which AS links lost routes and which AS links gained routes from the router's perspective, and looking at this graph can give a visual picture of the routing change event from one router's perspective.

A *Rank-Change* graph is weighted directed graph $G_{\Delta t} = (V, E)$ that captures the difference between a Link-Rank graph from time t_0 and a Link-Rank graph from time t_1 . The weight associated with each link indicates the change in the weight of that link. Thus, a positive weight would indicate a gain of routes carried over the link, while a negative weight would indicate a loss. For ease of presentation, a *solid* directed edge corresponds to a rank increase and a *dotted* directed edge corresponds to a rank decrease.

More formally, let $G_0 = (V_0, E_0)$ and $G_1 = (V_1, E_1)$ be two Link-Rank graphs obtained from the same path vector router at times t_0 and t_1 respectively. Algorithm 1 constructs the corresponding Rank-Change graph:

Algorithm 1: Constructing A Rank-Change Graph

Input: Link-Rank graphs $G_0 = (V_0, E_0)$ and $G_1 = (V_1, E_1)$

Output: A directed Rank-Change graph $G' = (V', E')$

Construct a graph $G' = (V', E')$ such that $V' = \{\}$, $E' = \{\}$;

for each edge $\langle v, w \rangle \in E_0 \cup E_1$ **containing weight** $W(\langle v, w \rangle)$ **do**

if $W_1(\langle v, w \rangle) - W_0(\langle v, w \rangle) \neq 0$ **then**

$V' = V' \cup \{v, w\}$;

$E' = E' \cup \{\langle v, w \rangle\}$;

$W'(\langle v, w \rangle) = W_1(\langle v, w \rangle) - W_0(\langle v, w \rangle)$;

The algorithm constructs the Rank-Change graph by adding edges that have changed in Link-Rank values, to build the edge set, and building the vertex set for vertices associated

with these edges. The Rank-Change graph does not contain the edges that didn't undergo any rank change and as a result, the set of nodes may be reduced in the Rank-Change graph compared to $G_0 \cup G_1$.

For example, assume at time t_1 , the Link-Rank graph in fig 1(a) changes to the Link-Rank graph in fig 1(b). This could be due to problems in the peering between AS-2 and AS-4, or some other policy reason. In most cases, we do not know which link failures triggered the route changes, and we use Rank-Change graphs to help narrow down on most likely candidates. The corresponding Rank-Change graph is shown in Figure 1(c). The solid rectangle in the Rank-Change graph identifies the point of observation. Including the point of observation is useful when we combine Rank-Change graphs viewed from different points. In the example above, the links (1, 3) and (3, 4) have increased in ranks by 1, by virtue of the change in path for prefix B from 1,2,4,6 to the new path 1,3,4,6. Similarly, the links (1, 2) and (2, 4) have reduced in ranks by 1 (weight of -1). Thus, Rank-Change graph visually indicates only sections have changed.

2.3 Combining RankChange Graphs from Different Vantage Points

Each Rank-Change graph captures the dynamics as observed from the particular router's perspective, and for n monitoring points, we would have at most n Rank-Change graphs. Thus, there is a need to combine these views to construct a global picture. Algorithm 2 details how we can combine different rank change graphs to understand collective BGP dynamics. Since combining n rank change graphs can increase the size of the global graph by a factor of n in the worst case, it is necessary to abstract out most important changes from each viewpoint. We thus reduce each Rank-Change graph into a reasonably small graph representing the most important changes in the graph. These subgraphs are then combined into one graph by performing a union on the set of nodes, and a multiset union on the set of edges. We draw multiple links if required for AS-AS peerings that appear in more than one graph to later identify commonality.

Algorithm 2: Combining Rank-Change Graphs

Input: Rank-Change Graphs $G\{V, E\}$ for a Set of Routers P for time interval ΔT

Output: Global Graph $G'_{\Delta T} = (V', E')$ combining multiple views

Initialization: $V' = \{\}$, $E' = \{\}$;

for each $p \in P$ with Graph $G_p(V, E)$ **do**

Find a subgraph $g_p(v_p, e_p)$ of $G_p(V, E)$ with a threshold t_{thresh} to filter the most significant changes;

Construct Graph $G'_{\Delta T}(V', E')$ from the obtained $g_p(v_p, e_p)$ s as follows::

$V' = V' \cup v_p$;

$E' = E' \cup_M e_p$ where \cup_M is a multiset union operator;

end

2.4 BGP Link-Rank and Rank-Change Graphs

Link-Rank and Rank-Change graphs are defined in terms of arbitrary path vector routing protocols, but we use these graphs to study BGP routing dynamics and provide some insight into BGP. The BGP routing protocol is used to exchange reachability information between *Autonomous Systems* (AS) and a BGP path consists of a sequence of Autonomous Systems (ASes)*.

Each AS is modeled as a single *node*, and the peering sessions between routers in different ASs are modeled as *edges* in the graphs. The Link-Rank graph shows *logical links* between Autonomous Systems. Each logical link in the model may consist of several physical connections between two Autonomous Systems in the Internet. A path vector protocol such as BGP hides the details of these physical connections and only indicates the sequence of ASes used to reach the destination. It is also important to note that the Link-Rank graph does not provide a complete map of the Internet's AS topology. Instead, the link graph shows the ASes and inter-AS connections that are used by a single router. There will be additional ASes that are not visible from the router's perspective, and thus these ASes do not occur in the router's Link-Rank graph. In other words, the Link-Rank graph is a proper subgraph of the Internet's AS topology and represents one view of the preferred paths selected by a router.

Since the Internet is growing in size and a typical backbone router currently has routes to over 120K prefixes, it becomes useful to identify operationally important ASes by the number of prefixes they carry, and AS to AS level usage of BGP routes provides valuable insights on how routing changes take place. The complete Link-Rank graph contains several thousand ASes so in order to concisely present the graph and highlight bulk changes, we often show only the links whose rank exceeds some threshold value. Note that the objective is to understand BGP dynamics in the aggregate and the threshold value can be adjusted to vary the amount the dynamics of interest.

2.5 Extracting Event Information from BGP logs

In the construction of a Rank-Change graph, we compare BGP routes from two snapshots in the time-interval Δt . If Rank-Change graphs were to be laid out on a real time basis, then one critical aspect of capturing event dynamics would be identifying event boundaries in the BGP log. When a predetermined interval is not known advance, we select an interval with the help of thresholding and waiting.

In our context, one candidate for threshold parameter is the *magnitude of Rank-Change*. In a simplistic scheme, the threshold value T is pre-defined, and we start with a snapshot of the routes at time t_0 . Whenever the total Rank-Change crosses this pre-defined threshold T , we start a convergence timer (around 5 mins) to reduce the impact of BGP convergence. When the timer expires at time t_1 , another snapshot of the system is taken, and with these two snapshots, we can construct the Rank-Change graphs for $\Delta t = t_1 - t_0$. For the next sequence, t_1 acts as the starting point and Rank-Change graphs are constructed for $\Delta t = t_2 - t_1$. Further improvements on this basic threshold

*AS path attribute may include both AS sequences and AS sets that result from aggregation, but AS sets rarely occurred in the data we observed




Figure 2 Effects of different *threshold* values on identification of number of events for January 2003

scheme may be achieved, by employing other feedback based learning techniques, like the *beta-gamma* adaptive threshold regulation method explained in [8]. Figure 2 shows the variation of the number of events captured with varying values of thresholds. As expected, increasing the threshold decreases the number of events captured.


3. Analysing BGP Path Dynamics

In this section we use the concepts from section 2 on BGP data collected from Oregon RouteViews[3] and RIPE NCC[2] over several months in 2002 and 2003. To illustrate the usefulness of the Link-Rank tools, three events are examined in detail. We pick these cases, out of the many results we had, to be representative of three different classes of events. All these events had similar update messages spikes, yet completely contrasting routing behaviors as seen by our graphs. In all three cases, we only show Link-Rank edges with a weight greater than 200. This value called *presentation threshold*, helps provide clarity in the presenting the graphs and has the effect of filtering out small changes. Presentation thresholds can be changed to vary the amount of details, according to the interest of the viewer. The greater the presentation threshold, the lesser the number of edges in the graph. We consistently maintain the same threshold of 200 for all our cases to establish a common ground for comparison.


The data retrieved from Ripe and RouteViews is first preprocessed to remove monitoring artifacts in the BGP data collection methodology[4]. In particular, if the BGP session between the monitoring point and the router being monitored breaks, then the entire routing table is exchanged after the session is re-established. This can result in large volumes of BGP updates, but only reflects re-establishment of peering session with the monitoring point.

3.1 Case I: 11th August 2003

As a first case, we examine the BGP dynamics on 11th August 2003, a date where no specific BGP related problems or anomalies were reported on standard operational mailing lists. Yet figure 3(a) shows a huge spike in the BGP updates on that day compared to other days. However, even for this simple example there is a large amount of logged data and understanding the activity during even this “normal” time period can be a challenge for operations. Figure 3(b) shows the Rank-change graph around 7:36 GMT, and the edge



(a) Number of Updates in 6 hour bins from Aug 11th to Aug 13th 2003



(b) Rank-Change Graph for AS 3130 on Aug 11th 2003

Figure 3 Case-I

weight for each link counts the change in Link-Rank in the last 5 mins. Note that all the changes are observed from vantage point connected to AS 3130. Rank-Change graphs from other monitoring points do not show any significant rank changes. Update logs confirm that in the interval between 7:25 and 7:40, only one AS of the over 30 monitored AS, shows high routing activity corresponding to 3(b).

The graph shows that around 10,000 prefixes stopped using a path beginning with $\langle AS3130, AS1239 \rangle$, and most of them switched to a path beginning with $\langle AS3130, AS2914 \rangle$. Below AS 1239, the path changes are spread across multiple links to 4323, 4637, 7132, 7018 and 3356. But large scale path changes did not occur below the first two hops. The Rank-Change graph for a single AS achieves the first step of compactly representing a high number of updates to give a simple picture of how the routes from a single router's viewpoint changed. Note that in this case, the other routers do not see any significant rank-changes and hence the combined graph from all routers is simply the Rank-Change graph from figure 3(b).

Although the Rank-Change graph shows only two neighbors for AS 3130, there are many more neighbors present in the update logs that are omitted due to rank changes less than presentation threshold(200 in this case). Looking at the distribution of rank changes lower down, it seems likely the routing changes observed are due to some local peering problems between AS 3130 and AS 1239 or some internal problems at AS 1239 or AS 3130. As observed from all other points, none of the links underwent a loss or gain of more than 200 prefixes. Rank-Change graphs for later times during this day, produced more such large routing swings to and away from AS 1239, observed from the same AS. Through this quick and simple analysis, and examining continuous Rank-Change graphs for extended periods, an operator can quickly identify the dominant BGP dynamics during this time period.

Note that in Figure 3(b) (and other later Rank-Change graphs), we do not see a perfect match in the number of prefixes lost by an AS on one link, and the number of prefixes gained by that AS on other links. In this case, AS3130 loses 9871 prefixes on the peering with AS 1239 and there is a gain of 8771 prefixes on the peering with AS 2914. The difference can be due to the slight variation in the total number of reachable prefixes at any instant of time. The presentation threshold can also be responsible for filtering out some of the smaller changes.

3.2 Case II: October 3, 2002

In a second example, the Link-Rank tools detected some unusual activity around 12:10 pm GMT on October 3, 2002. The algorithm produced several Rank-Change graphs from various monitoring point during that day and showed large link rank changes from multiple points. Figure 5(a) shows the number of updates in 6-hour intervals around October 3, 2002 with a huge spike.

Figure 4(a) shows the Rank-Change graph for AS 6539. Note that although, the view is observed from AS 6539, none of the links adjacent to AS 6539 actually appear in the graph. This is because none of these direct links show significant changes. The figure shows 6506 routes switched from AS 701 to other ASs. Similar to our analysis in the Case 1 example, an operator may suspect the source of the event is either AS 701, or the peering between AS2914 and AS701. But unlike Case 1, this event also generated Rank-Change graphs for many other routers being monitored and Figure 4(b) and Figure 4(c) show the Rank-Change graph as viewed from two other ASs, AS 3130 and AS 2914 respectively.

To better understand the event, we can combine Rank-Change graphs for different observation points into a single representative graph. We use algorithm 2 and excluding edges with changes less than 1000. Figure 4(d) shows the resulting combined view. We can see a high number of dashed edges into AS 701, as well as no solid edges to AS 701. The high in-degree of the dashed edges in the global graph suggests that AS 701 lost routes observed from various vantage points and is the likely cause of event in this case. This also indicates this event had a significant impact on all of our monitoring points. Furthermore, we can even identify common segments, such as the peering between AS2914 and AS701, that lost routes as viewed from multiple points. Our analysis is confirmed by reports about traffic congestion with AS 701 on the North American Network Operators (NANOG) mailing list[9]. We can see that the Link-Rank tools not only help an operator identify the potential problem, they provide a concise description of large volumes of data and allow the operator to get a visual feel of the BGP dynamics resulting from this change.

4. BGP Dynamics During the SQL Worm

As a final representative case, we look at the BGP activities during the SQL Slammer worm attack. The SQL Slammer worm alias "Sapphire" infected around 74855 hosts in a period of 30 minutes between 5:29 GMT and 6:00 GMT [10] in January 2003. Beginning around 05:00 Sat Jan 25 2003 GMT, the MS-SQL Server Worm began propagating through the Internet. Infected hosts generated a high volume of malicious worm probe traffic to UDP port 1434. Traffic congestion led to denial of service against many hosts and as well as

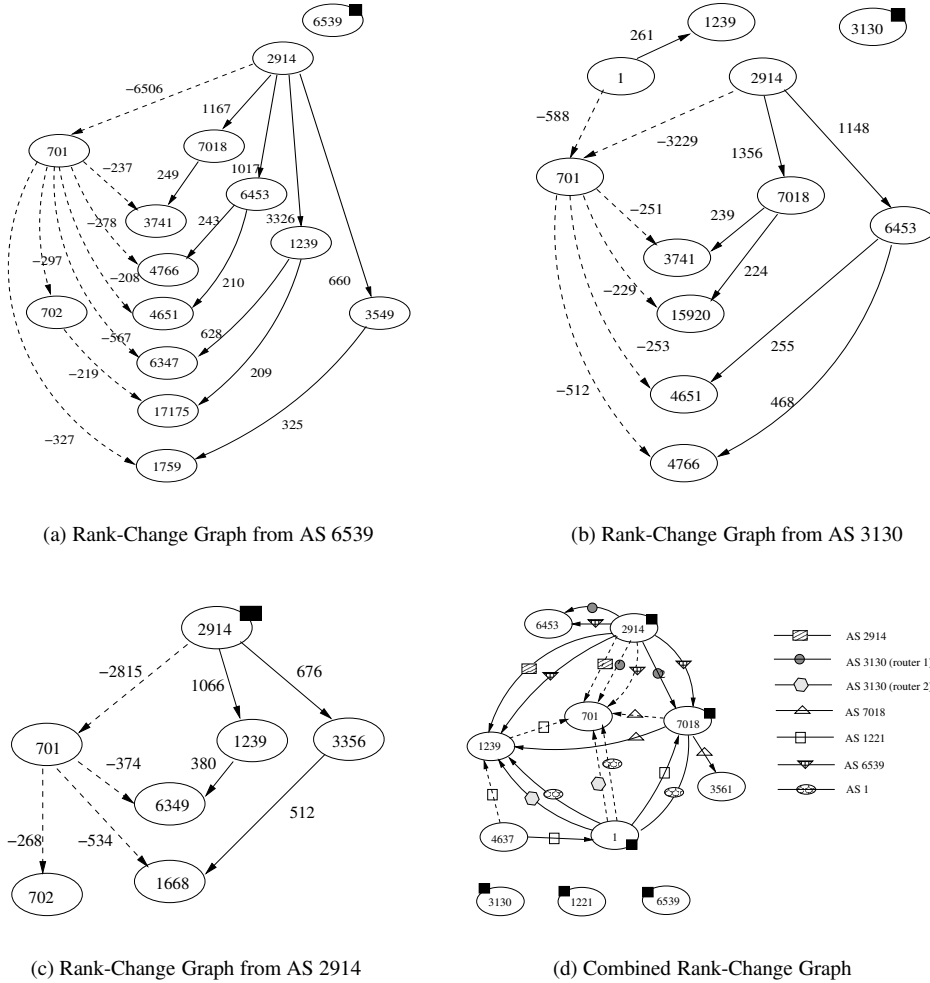




Figure 4 Case II: October 3, 2002

the network infrastructure. Concerns were raised on how Internet Routing was affected by the worm and [11] showed that around 5,000 prefixes were instantly withdrawn after the worm started infecting hosts. Internet health report for that time [12] showed that the delays on many links were above the critical value. Though Sapphire did not carry any malicious payload, its damage was caused by overloading networks by saturating the bandwidth and taking database servers out of operation. As a result, many individual sites reportedly lost connectivity as their access bandwidth was saturated by local copies of the worm. To analyze the amount of Internet backbone disruption, we use Link-Rank to examine the extremely large volume of BGP data collected during this time period. Figure



(a) Case-II, October 2002



(b) Case-III, January 2003

Figure 5 Update Count Plots

5(b) shows the number of updates observed at all the monitoring points at RouteViews in 6 hour bins. Once again, we can see the spike in the updates around the worm attack.

Using the concepts from section 2, we constructed Rank-Change graphs from various vantage points during the time window 05:25 am to 05:37 am GMT on Jan 25, 2003. To maintain consistency, in figures 6(a), 6(b) and 6(c), we show the Rank-Change graphs for the same three monitored AS used in Case-I and Case-II. We see that in all three cases, the dashed lines (route loss) dominate. This suggests that routes lost on one sub-path are not compensated by a gain on another. In other words, the routes were withdrawn along the shown AS-AS peerings, but for some reason not transferred to another AS-AS peering. Dashed lines (or Route loss links) dominating the Rank-Change graphs is not very common, and is an indication of loss of reachability. Whether the reachability is affected by problems at the edge, core or close to the monitoring point, may not necessarily be understood with the help of Rank-Change graphs and would require operator expertise.

Figure 6(d) shows the combined graph for these three monitoring points. A main result is actually the sparse nature of both individual and combined Rank-Change graphs. Although the volume of BGP updates was very high at each of the monitoring points, there were no large-scale common changes. This suggests that the changes were distributed widely with no one link suffering a large hit. The analysis on SQL slammer worm in [10] shows that a large number edge ASs were affected and the damage caused was geographically diverse.

We showed the effectiveness of Link-Rank using three diverse cases of BGP routing changes. In the first case, although a large number of route changes were observed at one point in the network, such large-scale change did not show up on the Rank-Change graph drawn from other vantage points, indicating local dynamics that have little impact on the rest of the Internet. In the second case, although the amount of route changes observed from the individual points look moderate, the Rank-Change graphs from multiple points reveal a common pattern of routes moving away from the same AS, indicating that the event had a much wider impact on the overall Internet traffic flow. In the third case, an exceptionally high volume of BGP updates was observed during the worm attack on Jan 25th, 2003. However, the Rank-Change graphs showed that the magnitude of changes

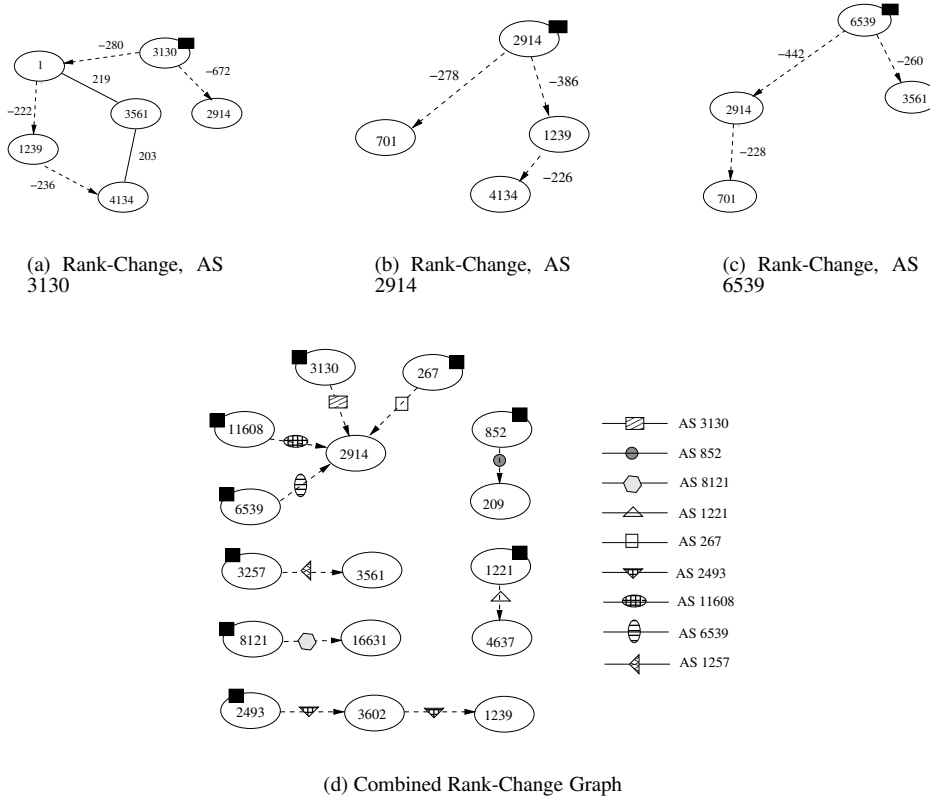


Figure 6 Case III: Slammer Worm attack

were much less and there were hardly any common pattern of route changes. Changes occurred at different parts of the Internet, indicating the distributed nature of the event.

Note that although we used Link-Rank in constructing Rank-Change graphs, for brevity, we did not explicitly present Link-Rank graphs in our case case studies. While selecting views to combine for a better picture, the LinkRank graph is very useful here to see how important the peerings are and how topology pictures compare in order to select peer routers that contribute useful and significant information about the Autonomous Systems and/or links of interest.

5. Related Work

In this paper, we analyzed AS path information from BGP update messages. [13] and [14] also used AS path information from BGP routing tables, but their objective was to map and derive characteristics of the Internet inter-domain topology. [15] further inferred the relationships between ASes using AS path information. [16] also inferred inter-AS relationship, but their approach is based on the consistency among multiple routing tables.

These studies have led to a better understanding of the structure of the Internet topology and the commercial relationship between ASes. However, each of them gives us only a snapshot of the Internet and tells us little about how the Internet routing changes. Our study, on the other hand, follows AS path changes and their effects on the overall routing topology to reveal the source and impact of routing problems. Researchers have looked at BGP updates during stressful events such as the code red worm [17]. According to the study in [4], the worm had a big impact on some edge networks, and weaknesses in BGP's design and implementation substantially amplified the impact. Although this kind of studies offer some insights into the causes of the BGP updates in the routing events, they lack a formal model to systematically identify these causes and assess their impacts. This paper aims to provide such a formal model.

Projects like BGPlay [18] for visualization of inter-AS routing instability and HERMES [19] for visualization of inter-AS connectivities are examples of successful BGP visualization. [20] also describes a good visualization tool for BGP called Elisha. Though these visualizations are related, none of these tools directly capture BGP dynamics in the sense discussed in this paper.

6. Conclusion

The large volume of BGP log data obscures the view of BGP routing dynamics and makes it difficult to extract significant routing outages from routine updates. Furthermore, the limited view of individual vantage points from which the data is collected also severely constrains the ability to derive the causes and assess the impact of observed routing dynamics. As a first step towards a global Internet routing monitoring toolset, the Link-Rank design explored a new approach for presenting the routing changes in a concise graphic form. Simply weighting the links individually based on prefixes reached per vantage point, Rank-Change graphs were able to capture AS level dynamics and give insight about the event. Combining Rank-Change graphs from multiple points allowed us to narrow down on the likely cause of change.

We believe Link-Rank can be used as an effective monitoring toolset either at a single monitoring point to draw a Rank-Change graph, or at a BGP monitoring site collecting updates from multiple vantage points. A network operator can use the former to monitor the paths used by his own network, and contrast the latter against the locally observed BGP events to assess the scope of the impact. As shown in our case studies, these graphs can be used to pinpoint the cause of change and evaluate its impact on the rest of the Internet. Our future work includes the use of IP address space metric in combination with the prefix count metric as the link weight in drawing routing graphs, the consideration of the tier level of ASes in the rank change analysis, and further improvements in combining multiple views to scale well with increasing numbers of vantage points.

References

- [1] Y. Rekhter and T. Li, "A border gateway protocol (BGP-4)," *Request for Comment (RFC): 1771*, Mar. 1995.

- [2] RIPE, “Routing Information Service Project,” <http://www.ripe.net/ripence/pub-services/np/ris-index.html>.
- [3] Univeristy of Oregon, “The Route Views Project,” <http://www.antc.uoregon.edu/route-views/>.
- [4] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. Wu, and L. Zhang, “Observation and analysis of BGP behavior under stress,” in *Proceedings of the ACM SIGCOMM Internet Measurement Workshop 2002*, Nov. 2002.
- [5] R. Govindan, C. Alaettinoglu, G. Eddy, D. Kessens, S. Kumar, and W. Lee, “An architecture for stable, analyzable internet routing,” *IEEE Network*, January/February 1999.
- [6] Ramesh Govindan and Hongsuda Tangmunarunkit, “Heuristics for internet map discovery,” in *IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000, IEEE, pp. 1371–1380.
- [7] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, “On inferring as-level connectivity from bgp routing tables,” Technical Report UM-CSE-TR-454-02, 2002. <http://topology.eecs.umich.edu/>.
- [8] Chengxiang Zhai, Peter Jansen, Emilia Stoica, Norbert Grot, and David A. Evans, “Threshold calibration in CLARIT adaptive filtering,” in *Text REtrieval Conference*, 1998, pp. 96–103.
- [9] “NANOG mailing list,” <http://www.nanog.org/maillinglist.html>.
- [10] David Moore et. al., “The spread of the Sapphire/Slammer worm,” <http://www.cs.berkeley.edu/nweaver/sapphire/>.
- [11] Tim Griffin, “BGP Impact of SQL Worm,” http://www.research.att.com/griffin/bgp_monitor/sql_worm.html.
- [12] Internet Health Report, “Sapphire Worm Attack,” http://www.digitaloffense.net/worms/mssql_udp_worm/internet_health.jpg.
- [13] R. Govindan and A. Reddy, “An analysis of inter-domain topology and routing stability,” in *Proceedings of the IEEE INFOCOM '97*, Apr. 1997.
- [14] M. Faloutsos, P. Faloutsos, and C. Faloutsos, “On power-law relationships of the internet topology,” in *Proceedings of the ACM SIGCOMM '99*, 1999.
- [15] L. Gao, “On inferring autonomous system relationships in the internet,” *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 733–745, 2001.
- [16] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, “Characterizing the internet hierarchy from multiple vantage points,” in *Proceedings of the IEEE INFOCOM '02*, New York, NY, June 2002.
- [17] CERT Advisory CA-2001-19, ““Code Red” Worm Exploiting Buffer Overflow In IIS Indexing Service DLL,” <http://www.cert.org/advisories/CA-2001-19.html>.
- [18] BGPlay, “BGPlay,” <http://www.dia.uniroma3.it/compunet/bgplay/>.
- [19] HERMES, “HERMES Visualization of Internet Service Provider Relationships,” <http://www.dia.uniroma3.it/hermes/>.
- [20] ELISHA, “ELISHA Visualizing the Internet Anomalies and Dynamics,” <http://www.ripe.net/ripence/pub-services/np/ris-index.html>.