

Visualizing Internet Routing Dynamics using Link-Rank

Mohit Lad, *Member, IEEE*, Dan Massey, *Member, IEEE* and Lixia Zhang, *Senior Member, IEEE*

Abstract—Visualizing Internet routing is an important first step for understanding the complex dynamics involved. Given gigabytes of routing data, extracting the nature and type of events from this data is a major challenge due to the size and multi-dimensionality of the data. In this paper, we present Link-Rank, a tool that visualizes Internet routing changes. Our approach captures information about links that lost routes and links that gained routes in a compact form, thus enabling one to summarize huge amounts of routing data and gain insight into network events. Using case studies, we show how Link-Rank is useful to provide insight into network routing dynamics.

Keywords

Inter-domain Routing, Route Visualization, Routing Dynamics.

I. INTRODUCTION

Routing protocols in a network are responsible for establishing routes to different destinations in the Internet. In recent times, a lot of interest has been generated on understanding and modeling the behavior of Internet routing. A major challenge for visualizing the dynamics of Internet routing, is the sheer size of the Internet routing data to be analyzed. A single router in the Internet could get over 30 Mega bytes of uncompressed routing data from each neighbor per day. Thus, a router connected to about 15 neighbors could receive could receive over 0.5 Gbytes of data per day. Further studies have shown that this amount of data increases multifold times during events of network stress like worm attacks [18].

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No DABT63-00-C-1027 and by National Science Foundation(NSF) under Contract No ANI-0221453. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors

Mohit Lad and Lixia Zhang are with University of California, Los Angeles, {mohit,lixia}@cs.ucla.edu, Dan Massey is with Colorado State University, massey@cs.colostate.edu

With the amount of routing data per day from a single router in the Internet being so much, visualizing routing as a whole by collecting data from multiple points becomes a very challenging problem. Besides the size of data, the multi-dimensional nature of the routing data makes the visualization problem even harder. Broadly, the data can be broken down into the three dimensions of temporal, spatial topological and spatial view. The temporal dimension involves routing events occurring continuously and organized by time. The spatial topological dimension concerns the presence of various nodes in the Internet each of which has some activity of interest. Finally, different observation points in the Internet see different routing behaviors based on events and their scope and the third dimension of spatial view concerns which observation point is collecting the data. While, this dimensionality makes visualization non-trivial, ignoring it can result in an incorrect or inaccurate understanding of routing. Thus, not only must a good visualization tool scale overall, but it must also scale along the individual dimensions of time, space and view.

In this paper, we present Link-Rank, a tool to visualize Internet routing dynamics. We take into account the multi-dimensional nature of the routing data and design Link-Rank as a wholistic solution while still addressing the challenge each dimension presents. Link-Rank uses a novel graph called rank-change graph that can capture changes in routes. We show using case studies how mega bytes of data can be easily summarized by these rank-change graphs. Link-Rank tool also contains a high level picture of routing dynamics called ‘activity plot’. With an activity plot, one can easily spot time periods of high routing dynamics without needing to look through lot of data. Often the needs of visualization would require one to be able to scale up or down on one or more of the three dimensions mentioned above, and Link-Rank allows users to vary the amount of detail they wish to see on each of these dimensions. Besides being useful as an operational tool, Link-Rank can also serve as a research tool to understand and analyze the behavior of Internet routing and evaluate the routing

protocol as a whole.

The remainder of this paper is organized as follows. Section II reviews Internet Routing and explains the data methodology. Section III introduces Link-Rank graphs with some simple examples. This section also introduces other features of Link-Rank tool used to deal with various challenges each dimension of data presents. Section IV uses real examples to show the usefulness of Link-Rank. This section also shows how views from different autonomous systems can be assembled into one routing view to better understand the dynamics. Section V discusses related work in this area, and finally, Section VII concludes the paper.

II. INTERNET ROUTING INFRASTRUCTURE AND BORDER GATEWAY PROTOCOL

The Internet consists of a large number of administrative networks called Autonomous Systems (AS). An AS may advertise one or more IP prefixes to its immediate neighbors. These neighbors, may in-turn propagate their routes to these address prefixes to other neighboring autonomous systems. BGP (Border Gateway Protocol [16]) is the routing protocol used by these ASes to exchange inter-domain routing information.

BGP routers propagate routing information with the use of ‘update’ messages. Each update message consists of one or more address prefixes and some path attributes. One of the most important path attributes is the *AS Path* that includes a series of AS numbers identifying the Autonomous Systems on the route to an address prefix. Each update is either an *announcement* to advertise/change a route, or it is a *withdrawal* to remove a previously announced route. In the latter case, the BGP update will not contain any AS Path attribute.

In BGP, two routers establish a peering session over a TCP connection. When the session first comes up, each router announces its entire routing table¹. After this initial route table announcement, the router should only send routing updates if a route changes. If the peering session breaks, all routing information learned from the peering session is discarded and the full table is re-announced when the session is re-established. As a result, one expects a large number of updates at the start of a peering session. Future updates should only occur when individual routes change.

Our study uses BGP data collected from well known BGP monitoring points. BGP monitoring projects at Routeviews[3] and RIPE [4] have set up BGP data

¹Routing policies are commonly used to limit the set of routes advertised to a peer or accepted from a peer.




Fig. 1. Internet routing and BGP monitoring

collection points in the Internet. These points act just like other BGP routers, except that they do not propagate any updates to the Internet. Figure 1 shows how a BGP path is propagated and also explains how data is collected at monitoring points. This figure also shows how a route to prefix P reaches the monitoring points. The monitoring point in this example connects to routers in both AS 7018 and AS 1239 and hence logs the BGP updates sent by these two autonomous systems. Any further changes to the route to reach P from either AS 7018 or AS 1239 will also be propagated to the monitoring point. The data collected by the monitoring points provides a valuable resource to investigate various issues in the operation of Internet routing. The data collected at these points contains the following key fields

$$(routerIP, type, prefix, newPath, attributes)$$

where routerIP is the IP address of the router in the AS the monitoring point connects to, type indicates whether the update is a withdraw and will not contain any new path or is an announce and will contain a new path, prefix contains the destination prefix the update is for, newPath is the path to be used for announcements, and attributes represent a few other optional attributes. For more details, the reader is encouraged to refer to [13]

III. LINKRANK AND RANKCHANGE GRAPHS

In this section, we introduce the LinkRank toolset and explain the various aspects of Link-Rank. One of the main challenges LinkRank deals with is processing huge amounts of BGP routing data. As mentioned before, the routing data volume can be broken down into the three dimensions of topological granularity, time and view. For a visualization tool to be able to deal with this volume of data, the following questions need to be addressed.

- 1) Topological Granularity: How to decide which links and autonomous systems to visualize out of thousands?
- 2) Time: How to pick time intervals of high routing dynamics?
- 3) View: How to combine views from multiple observation points?

We start by explaining how to deal with high volume of data due to topological granularity.

A. LinkRank Graph

The objective of a Link-Rank graph is to represent the links between autonomous systems from a particular view point, in a visually compact form. Often, it is difficult for an autonomous system to know which links beyond their neighbor are heavily used to reach other destinations in the Internet. With LinkRank graphs, it is easy to spot the highly ranked edges which carry a high number of routes to other destinations.

We start with a simple case of a single view point M , and attempt to visualize routes for a given time instant t . The challenge we face now is to control the level of detail of routes from a single view point. The first graph in Figure 2 consisting of nodes 1, 2 and 3 shows a Link-Rank graph at some time t . Here the link weights are the routes reached by 1 through that link. Note the direction of an edge indicates the direction a data packet would travel when reaching these destinations. Each link in the Link-Rank graph is assigned a weight based on the *importance* of the link. One measure of importance used in this study is number of routes that rely on that link. More precisely, $rank(\langle v, w \rangle \in E) =$ the number of paths, $path_i$, that contain subsequence $\langle v, w \rangle$.

B. Rank-Change Graphs: Tracing Link-Rank Changes Over Time

A Link-Rank graph from a point provides a snapshot of the autonomous system links used to carry routes. An update message received at a BGP router, may change the path to some destination. Thus, every update received at that point of observation could change the weight of one or more links in the LinkRank graph. Figure 2 shows a sequence of updates causing routes to move from the link (2, 4) to link (2, 7). In this case, every update reduces the rank of link (2, 4) by 1 and increases the rank of link (2, 7) by 1. Given this continuous change in the state of the LinkRank graph, how does one decide at what stage should one view the LinkRank graphs. In other words, how does one capture the change in routes. We deal with this




Fig. 3. Rank-Change graph from two Link-Rank graphs

challenge by using the idea of a ‘change threshold’. A snapshot of LinkRank graph is taken only if the running total of rank of any link changes by more than the ‘change threshold’. In Figure 2, a snapshot is taken with the change threshold value set to 100. The change threshold allows a user to decide what magnitude of changes the user want to catch. By using the change threshold, one can condense a series of updates into two LinkRank graphs. Even in this case, there may be some links whose weights have not significantly changed or changed at all. For this purpose, we define Rank-Change graphs, which capture the sections of the routes that change. Conceptually, a Rank-Change graph denotes the change in BGP routes as observed from an observation point over a period Δt . The Rank-Change graph shows which links lost routes and which 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. Figure 3 shows the starting and ending LinkRank graphs and also the corresponding RankChange graph. The Rank-Change graphs are an important first step in visually analysing the BGP update dynamics.

More formally, A *Rank-Change* graph is a 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 *green* directed edge corresponds to a rank increase and a *red* 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:

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




Fig. 2. Changes in Link-Rank due to update messages

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**

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

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

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

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 $G_0 \cup G_1$. 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 autonomous system links by the number of prefixes they carry, and this level usage of BGP routes provides valuable insights on of how routing changes take place. The complete Link-Rank graph contains

several thousand autonomous systems. The way we construct Rank-Change graphs, allows us to condense these graphs into ones with lesser nodes and edges. Further, the concept of time window presented later allows us to group events based on the time granularity a user is interested in.

C. Activity Plots: Scaling on Time and View Dimension

From the BGP data, we found the amount of routing activity to differ from time to time. Since there is no index capturing this amount of routing activity, it becomes difficult to differentiate periods of high routing activity from low activity. As a result, it becomes necessary to go through a huge amount of data, when if some index of routing activity was present, one could look through much lesser data. In addition, one could easily characterize different observation points for the amount of activity. With this objective, we added the activity plot feature to the Link-Rank tool. This activity plot provides a high level picture of routing dynamics in a time period, and helps in identifying periods of high routing dynamics. Figure 4, where X-axis is time, shows an activity plot for a week as observed from one observation point. Each activity bar indicates the magnitude of rank change activity during that time. The absence of activity bars during other times in Figure 4 indicates lack of activity exceeding the user defined

threshold of route changes. The activity plot can be thought of as a plot of routing ‘health’, with clustered bars indicating instability in routing, while no bars indicating changes below acceptable level.

Figure 5 shows how we plot an activity bar for a given time based from a rank change graph. In the example, the total gain is 200 and the total loss is 100. Hence the green bar is longer than the red bar. Thus activity plots could be a good indication of kind of network dynamics. A higher gain (green) than loss (red) could be due to a combination of longer new paths as in Figure 5 and new routes being announced which were previously unavailable. On the other hand, a series of red activity bars without any green bars could imply a loss in connectivity for some prefixes. Activity plots are very useful in investigating network events and can be used as a starting point to pick time intervals of high routing dynamics.

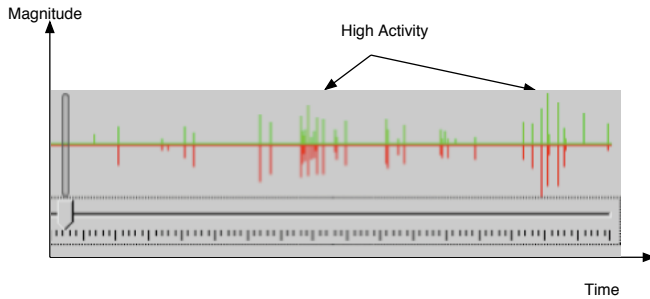


Fig. 4. Activity plot for a week from one view




Fig. 5. Plotting an activity bar

Besides, the time dimension, activity plots can also help identify views with common behavior. For example, Figure 6 shows activity plots from multiple views during the same time period. The highlighted area shows common activity behavior for these views at around the same time. Thus, when augmenting a view with other views to narrow down on the cause, one can look at activity graph to identify potential views

which show behavior at around the same time. We use an example later to show how this can be done.




Fig. 6. Activity plots from different views

D. Time Windows and Drilling Down

Another useful feature of LinkRank toolset is its ability to change the longevity of rank changes. Due to the presence of slow convergence [12], we often see some invalid paths and these could appear as genuine route changes which are short lived. With the time-window control, one can increase/decrease the required longevity of events in order to be interesting to the user. For instance setting the time window to 15 minutes will ensure that events lasting more than 15 minutes will not be missed, no matter where we position the time window. Thus by setting a time window to δ , the user is willing to overlook events with a durations shorter than δ . However, it is not necessarily the case that events of a shorter duration are not captured. If the start of the event is close to the end of the time window, then the event will show up. In other words, by setting a time window of δ , the tool will never miss any event lasting longer than δ , but may or may not show events lasting shorter than δ .

At the same time, LinkRank provides a visual hint in the form of activity plot at the bottom on high routing activity. The first part of Figure 7 shows a case where there is not much change indicated in the rank-change graph, but as can be seen the time-window contains a

lot of activity. To examine this activity in more detail, one could either reduce the time-window, or use the drill-down feature of the tool. The drill-down feature creates a new window and expands the activity inside the current time-window to a larger time-span so as to go into greater details. In Figure 7, we can see a lot of activity in the initial half of the expanded plot. Further examination revealed continuous fluctuations of thousands of routes between two links causing this disturbance.

Thus both time-window control and drill-down feature present an opportunity to increase or decrease the level of detail as per needs.

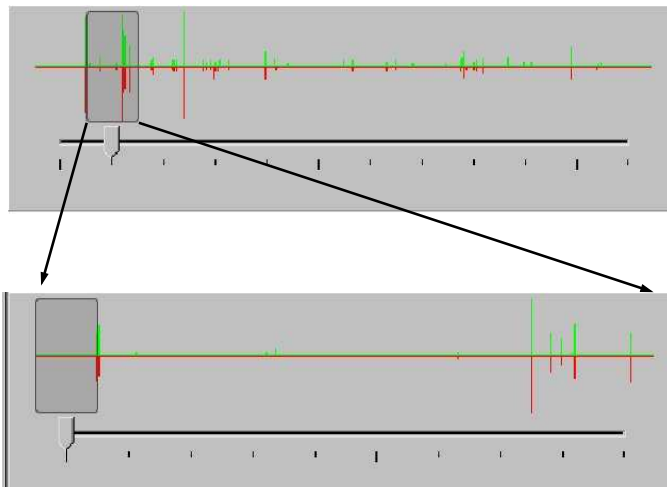


Fig. 7. Drilling down to increase level of detail in activity

IV. USING LINKRANK TO UNDERSTAND ROUTING CHANGES

LinkRank can be used for two types of discoveries. First, LinkRank can be used to show what happened in the Internet in a given time period. Here, LinkRank can summarize route changes and show which AS-AS links were changing. Further, the second type of discovery LinkRank can be used for is ‘isolating’ the cause of the observed routing changes. We use case studies in this section to show how LinkRank can be used for both kinds of discoveries.

A. Understanding routing dynamics

We explain how LinkRank can be used to understand what happened with the help of an example from real data from July 13, 2004. We observed a lot of routing activity from AS 286. Figure 8 shows the view from AS 286 during this period of high activity. From the activity graph, it can be seen that this graph falls in a period

of high activity. When we sequence through the time period with the help of LinkRank animations, it can be seen that during this period, the routes keep switching between the subpaths (286, 174) and (286, 209, 174). Each link in this graph has two weights, the first being the Link-Rank weight (number of absolute routes using the link) and the other is the change in Link-Rank weight (the gain/loss of routes). This time period generated thousands of updates for this view point, while using LinkRank, it is possible to summarize the routing changes in one line.

B. Identifying cause of change

Having seen how LinkRank can tell us the nature of routing dynamics, here we discuss how LinkRank can be used to perform diagnosis. Figure 9 shows the RankChange graphs during for a 10 minute interval around 20:15 GMT on July 14, 2004. We can see there is a loss of around 250 routes (for brevity we do not show the absolute rank values here) along the next hop of AS 2914 and the corresponding gain is of around 250 routes on the next hop of AS 1239. Note the blue color indicates the node is the observation point. In addition, the orange colored nodes indicate the nodes AS 2914 and AS 1239 are also observation points. This particular route change could have happened due to some problem on the red path, or due to some recovery of a previous problem on the green path. It is not possible to make conclusions based on just this graph, so we take advantage of the fact that we can also see the view from AS 2914 and AS 1239. Observing from AS 2914 during the same time period does not reveal dynamics involving any links common to that observed from AS 3130. Thus, it is unlikely there was a problem on the routing path via AS 2914. Next we examine the view from AS 1239. We can see from Figure 10 during the same period as in Figure 9, the path from AS 1239 to AS 8437 got shorter. As a result, in Figure 9, the path via (3130, 1239, 8437) got shorter. Thus, with lots of view points, one can expect to perform better diagnosis by using the views from the orange nodes on potential problem areas.

C. Assembling Multiple Views

So far we have considered visualizing a single observation point. In this section we motivate the need for assembling multiple views together and present some interesting cases to show how assembled views helps. Often routing changes observed from different monitoring points are related to each other. This relation cannot be easily observed when viewing from

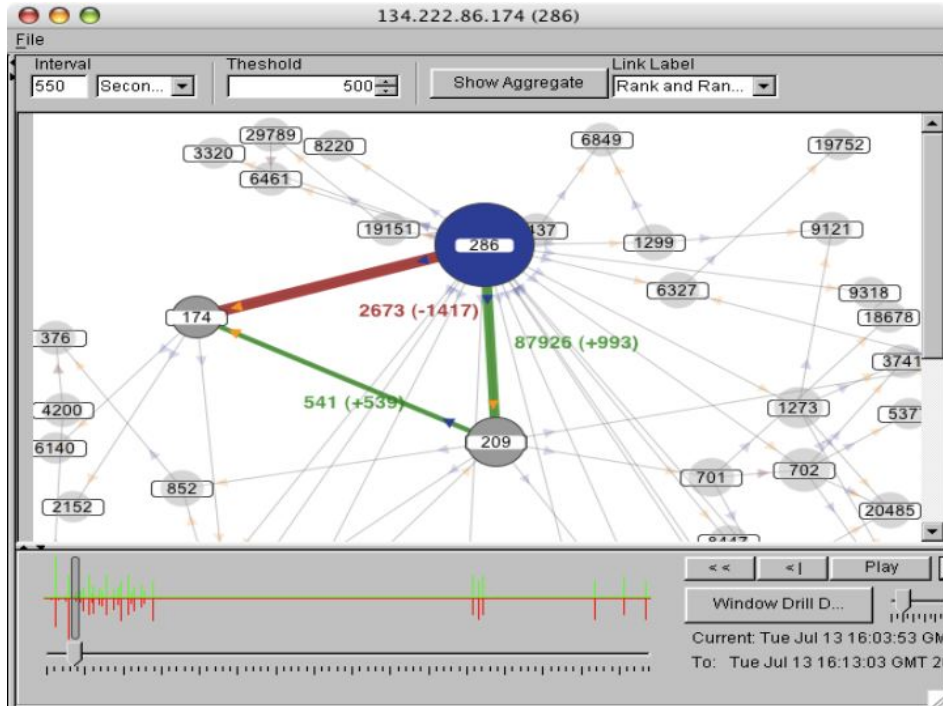


Fig. 8. View from AS 286

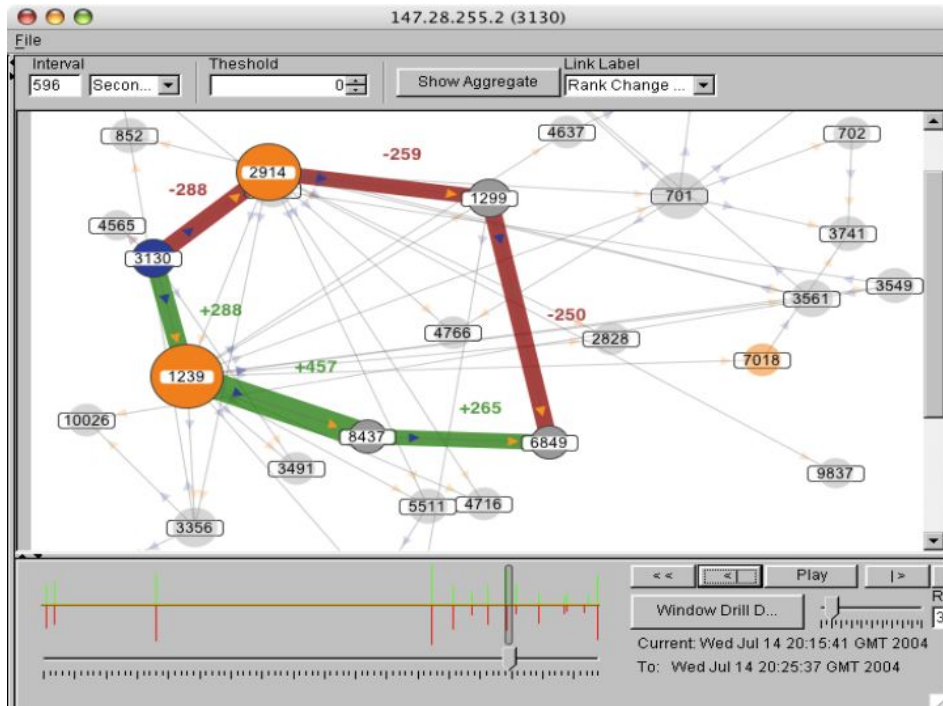


Fig. 9. View from AS 3130

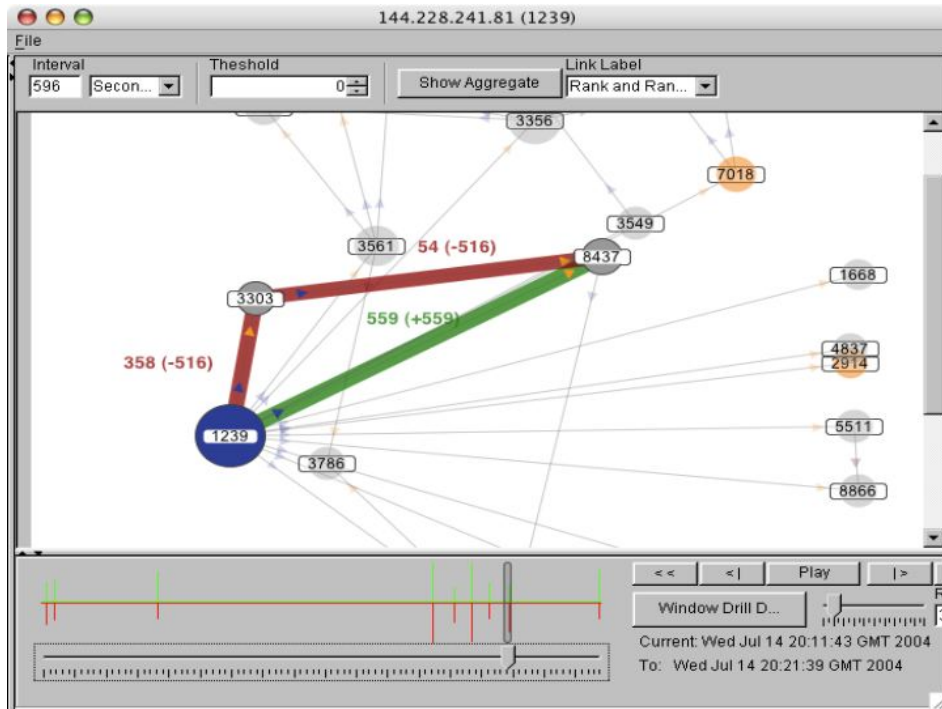


Fig. 10. View from AS 1239

a single point of view. The Link-Rank tool has the ability to assemble the views from multiple observation points in a single window. This enables the user to identify common behavior as well as better understand the triggering cause of updates received at multiple points. The tool distinguishes each observation point by assigning that point and all its related changes a single color. Since each observation point is color coded, we cannot use red and green for loss and gain in ranks and instead use dashed lines for rank loss and solid lines for rank gain. As before, circular nodes indicate the observation points to differentiate from other nodes. The observation points to be used for assembling views are picked by the user from the list available in the main window. A time is specified and as a result an assembled view is generated. The assembled view activity plot now contains a combined activity for all the views being assembled. All the timing controls used in the single point case can be used just as well here.

We present an example of how combining views can help understand network events better. Figure 11 shows an assembled view from two observation points at the same time. This view clearly shows common behavior observed at these two points. The link between AS 1239 and AS 20115 is seen to be the common factor in the route changes observed from both observation points.

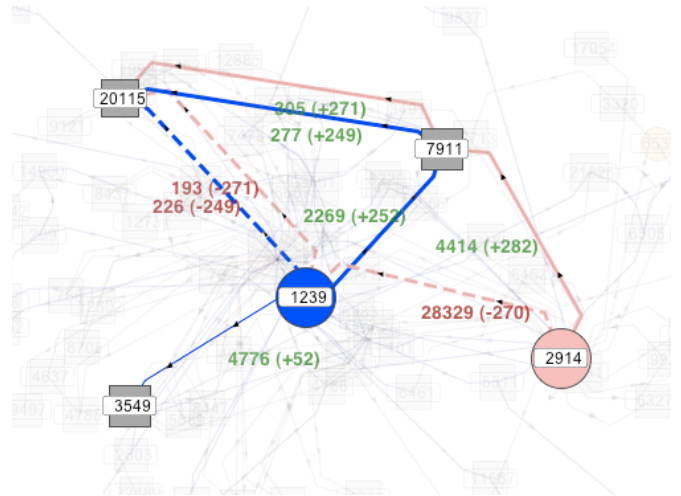


Fig. 11. Assembled views in Link-Rank

V. RELATED WORK

Visualization in the Internet has been proposed before for various problems. A popular visualization project in Internet is Caida's Internet Visualization [5]. They created a map for the Internet connectivity by first identifying routers and then associating them with autonomous systems. The map in [5] contains 12517 Autonomous Systems and 35,334 peering sessions. Each AS node is assigned a position based on polar co-ordinates links

are color coded to indicate the outdegree of an AS. This map provides some interesting insight into how connectivity varies for different geographic locations. However, the map provided by Caida is not updated since its release. A related project on AS Topology at UCLA [19] provides an updated AS level topology in the form of text files. A detailed explanation of their methodology and collection is explained in [20]. Another tool for visualizing inter-AS connectivity is HERMES [2] developed at University of Rome.

The Internet research community has also benefited from the presence of certain other tools that use BGP data to provide insight into operations in the Internet. The Security Visualization project [15] provides Elisha among other tools for visualizing various aspects of the Internet. ELISHA was designed for understanding Internet anomalies through visualization of BGP data. Another related project on visualizing BGP dynamics is BGPlay [1]. BGPlay is an interesting tool which shows using animations how routes change to a particular prefix. However, with over 130,000 prefixes in the Internet, this tool cannot be used for understanding BGP dynamics in the aggregate. We feel this tool complements the goal of Link-Rank and can be used in conjunction with Link-Rank to provide more diagnosis power using visualization. In addition to academically developed tools, commercial tools like Packet Design Inc.'s Route-Explorer are also available for performing analysis of IP routing.

Besides visualization projects, there has been work on Internet topology discovery and BGP dynamics. Some representative work in AS topology inference includes [11], [10], [17] and [8]. Research on inference of cause of change from observed BGP updates includes [9], [6] and [7].

VI. OBTAINING THE LINK-RANK TOOL

Link-Rank is a java based tool implemented using a modified JUNG library available under open source license. Link-Rank is available under a GNU open source license to enable users to modify code if necessary. An alpha version of the Link-Rank tool was released on Dec 31, 2004. As of February 19, 2005 Link-Rank website has logged around 12000 page hits and 450 tool downloads. The Link-Rank website [14] provides details about the tool, including a user manual and frequently asked questions. The website also contains updated activity plots from major observation points to provide an overview of amount of routing activity going on in the last 7 days.

The Link-Rank website also maintains updated data from the RouteViews Oregon collector. This data can be accessed and downloaded through the website and used for visualization. A future release planned for May 2005 will contain the back-end code for processing raw updates so operators can visualize their own data which is not publicly available. This release will also contain features for synchronizing data between the client and the server. We are also working on combining views from multiple observation points for better diagnosis of network problems. This step is an extension of assembling views for diagnosis purpose.

VII. SUMMARY AND 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.

REFERENCES

- [1] BGPlay visualization [online, cited 02/25/2004]. Available from: <http://www.dia.uniroma3.it/~compunet/bgplay/>.
- [2] HERMES Visualization of Internet Service Provider Relationships [online, cited 02/25/04]. Available from: <http://www.dia.uniroma3.it/~hermes/>.
- [3] RouteViews Routing Table Archive [online, cited 02/25/05]. Available from: <http://www.routeviews.org/>.
- [4] Routing Information Service Project [online, cited 03/11/05]. Available from: <http://www.ripe.net/>.

- [5] Visualizing Internet Topology at a Macroscopic Scale [online, cited 02/25/04]. Available from: http://www.caida.org/analysis/topology/as_core_network/.
- [6] M. Caesar, L. Subramanian, and R. Katz. Root cause analysis of Internet routing dynamics. Technical Report UCB/CSD-04-1302, U.C. Berkeley, november 2003. Available from: citeseer.nj.nec.com/562290.html.
- [7] D. Chang, R. Govindan, and J. Hiedemann. The temporal and topological characteristics of BGP path changes. In *ICNP*, november 2003.
- [8] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proceedings of the ACM SIGCOMM '99*, 1999.
- [9] A. FeldMann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs. Locating internet routing instabilities. In *Proceedings of Sigcomm*, September 2004.
- [10] L. Gao. On Inferring autonomous system relationships in the internet. In *IEEE/ACM Transactions on Networking*, volume 9, pages 733–745, 2001.
- [11] R. Govindan and A. Reddy. An analysis of inter-domain topology and routing stability in Congested Networks: Experimentation and analysis of BGP update surge during slammer attack. In *Proceedings of the IEEE INFOCOM '97*, 1997.
- [12] T. G. Griffin and G. T. Wilfong. An analysis of BGP convergence properties. In *Proceedings of SIGCOMM*, pages 277–288, Cambridge, MA, August 1999. Available from: citeseer.nj.nec.com/griffin99analysis.html.
- [13] B. Halabi and D. McPherson. *Internet Routing Architectures*. Cisco Press, 2nd edition, 2000.
- [14] M. Lad, D. Massey, and L. Zhang. Link-Rank BGP Visualization [online, cited 02/25/2004]. Available from: <http://linkrank.cs.ucla.edu/>.
- [15] K.-L. Ma, S. F. Wu, and S. T. Teoh. Security Visualization [online, cited 02/25/05]. Available from: <http://www.cs.ucdavis.edu/~ma/SecVis/>.
- [16] Y. Rekhter and T. Li. A Border Gateway Protocol (BGP-4). *Request for Comment (RFC): 1771*, 1995.
- [17] 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, 2002.
- [18] 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*, 2002.
- [19] B. Zhang, R. Liu, D. Massey, and L. Zhang. Internet topology. <http://irl.cs.ucla.edu/topology/>.
- [20] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the internet's as level topology. In *ACM Sigcomm Computer Communication Review*, 2005.