Title
A Visual Exploration Process for the Analysis of Internet Routing Data

Authors
Teoh, Soon Tee
Ma, Kwan-Liu
Wu, Felix S.

Publication Date
2003

Peer reviewed

A Visual Exploration Process for the Analysis of Internet Routing Data

Soon Tee Teoh* Kwan-Liu Ma* S. Felix Wu*
Department of Computer Science
University of California, Davis

Abstract

The Internet pervades many aspects of our lives and is becoming indispensable to critical functions in areas such as commerce, government, production and general information dissemination. To maintain the stability and efficiency of the Internet, every effort must be made to protect it against various forms of attacks, malicious uses, and errors. A key component in the Internet security effort is the routine examination of Internet routing data, which unfortunately can be too large and complicated to browse directly. We have developed an interactive visualization process which proves to be very effective for the analysis of Internet routing data. In this application paper, we show how each step in the visualization process helps direct the analysis and glean insights from the data. These insights include the discovery of patterns, detection of faults and abnormal events, understanding of event correlations, formation of causation hypotheses, and classification of anomalies. We also discuss lessons learned in our visual analysis study.

CR Categories: H.5.2 [Information Interfaces and Presentation]: User Interfaces—Graphical User Interfaces (GUI); I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction Techniques

Keywords: information visualization, text visualization, network visualization , internet stability, homeland security

1 Introduction

The Internet has become an integral part of modern society. People depend on the Internet for the latest news, communication, and information on various subjects. World-wide commerce, government, industrial production, and financial transactions are just some of the activities which have become increasingly dependent on the Internet. Despite its apparent pervasiveness, the Internet is still growing at a tremendous rate, reaching more people in more places every year.

Because of the world's reliance on the Internet, it is critical to ensure that the Internet is functioning properly. Unfortunately, the Internet is vulnerable to attacks and errors. For example, worm attacks have on occasions propagated throughout the entire Internet, and effectively shut it down for hours, causing billions of dollars of damage, and affected people in other unquantifiable ways.

Much effort is therefore needed to protect the Internet against attacks and errors. The Internet is a complex distributed system running on a very large number of nodes using various protocols. A major effort in the authoritative Internet groups, such as the Internet Engineering Task Force (IETF) [IETF n. d.], is to understand the dynamic behavior of the Internet. The study of the operational behavior of Internet is thus a fundamental and important part of the effort to make the Internet more stable, robust, and secure. In the past, monitoring and analyzing network behaviors has been mainly done by browsing the raw data or looking at some simple plots of statistical analysis results. More sophisticated visualization techniques, which have been shown to be very helpful to many real-world applications, should also be developed and adopted for the analysis of Internet data.

Past efforts using visualization to better understand routing and Internet connectivity include Otter [Huffaker et al. 1999] and other tools by CAIDA [CAIDA n. d.], as well as the Internet Mapping Project [Cheswick et al. 2000]. These visualization tools give useful pictures of the Internet topology and changes in reachability. More detailed visualization-aided packet-level analysis of Internet data include [Estrin et al. 1999] and [Girardin 1999]. [Teoh et al. 2002] discusses anomalies found in the visualization of Origin AS changes.

We have developed a suite of visualization techniques and formulated a sequence of steps utilizing them for improved understanding of Internet routing data. Each step corresponds to a module in the interactive visualization system we have built. Following these steps to analyze the routing data proves to be very effective in finding the most important information hidden in the raw data. This paper describes how each step in the proposed visualization process is used to find correlations in events, discover patterns, detect faults and anomalies, make hypotheses about the causes of events, and classify different events. We show how this knowledge contributes to the understanding and security of the Internet, and also discuss the lessons learned from our experience.

2 BGP and Internet Routing

We examine Internet routing history to discover traces left behind by configuration errors and malicious attacks. Analyzing Internet routing logs can also help determine the stability of the current routing system.

Each network within the Internet is identified by its *IP prefix*. An example of an IP prefix is 128.120.0.0/16, which means every host IP address in the network shares the same first 16 bits: 128.120. One or more networks within a single administrative domain is referred to as an *Autonomous System* (*AS*), and is assigned a unique AS number.

Between two ASes, Border Gateway Protocol (BGP) [Rekhter and Li 1995] is used to exchange network reachability information so that routers can eventually forward data packets to the correct destination. BGP routers exchange messages in the form of BGP *announcements*. A BGP announcement lists a particular IP prefix and the path of ASes used to reach that prefix. For example, the BGP route "128.120.0.0/16: (7,23,92)" means that the IP prefix 128.120.0.0/16 could be reached by first going to AS-7, then to AS-23, and finally to AS-92. AS-92 is also known as the *Origin AS* or the *source* of 128.120.0.0/16. A BGP announcement can

*{teoh,ma,wu}@cs.ucdavis.edu

also be an explicit withdraw. For example, is an AS announces "128.120.0.0/16: WD", it means that the prefix is no longer reachable from the AS.

In this way, the connectivity of the Internet is maintained by routers communicating with BGP. To study the stability, problems, behavior and vulnerability of the Internet, we obtained from the Oregon Route Views server [Myer n. d.] the archived daily BGP routing data for the years 2000 and 2001. The Oregon Route Views server is thus our *observation point*. Each archived announcement is listed with its timestamp, which is the time that the announcement was received at the observation point. Because of the importance of BGP, and the complexity due to the distributed nature of BGP across the routers in the Internet, BGP dynamics have been analyzed in the past and continue to be a subject of much on-going research. Some previous studies of BGP dynamics include [Gao and Rexford June 2000; Griffin and Wilfong Aug 1999; Griffin and Wilfong Mar 2000; Pei et al. June 2002].

3 A Visual Exploration Process

To understand the dynamic behavior of Internet routing using BGP path announcements, a process of visual-based exploration is defined. The process consists of five steps:

1. Examining the plot of the number of announcements per unit time to get a sense of the overall distribution of the BGP path announcements. A general picture of the distribution and clustering of announcements is necessary for the user to develop an idea of what is normal/abnormal.

2. Browsing route announcement data to classify different types of instability events. Once clusters have been identified in Step 1, the next natural step is to examine each cluster. From the route announcement visualization, the user is able to see similar patterns in some clusters. Classification is important because similar patterns often imply similar causes.

3. Visualizing path changes to analyze each individual node and link. The goal is to identify properties and problems of different nodes/links/paths. This analysis can also track down the cause of an instability event.

4. Comparing AS path announcements from different peers. This also helps determine which nodes/links were responsible for the instability event observed.

5. Visualizing the distribution of different types of instability events. In Step 1, the distribution over time of AS path announcements is visualized. After analysis in Step 2, each cluster observed in Step 1 is classified into a different type. The visualization in this step incorporates the results of the analysis in Step 2, thereby providing richer information.

The rest of this section describes the five modules in our system, and how they relate to one another. We also mention some routing problems discovered and insights gained.

3.1 Aggregate Data Browsing Module

Beginning with the broad goal of using visualization to analyze BGP AS route announcements for understanding the dynamic behavior of Internet routing, we asked the question "*Are there any patterns or clusters in the timing of the announcements*?" The Aggregate Data Browsing Module was thus designed to give a sense of the overall distribution of the announcements. Aggregate data visualization refers to the plot of the number of BGP path announcements and the unique routes used per unit period of time, as described in [Teoh et al. 2002].

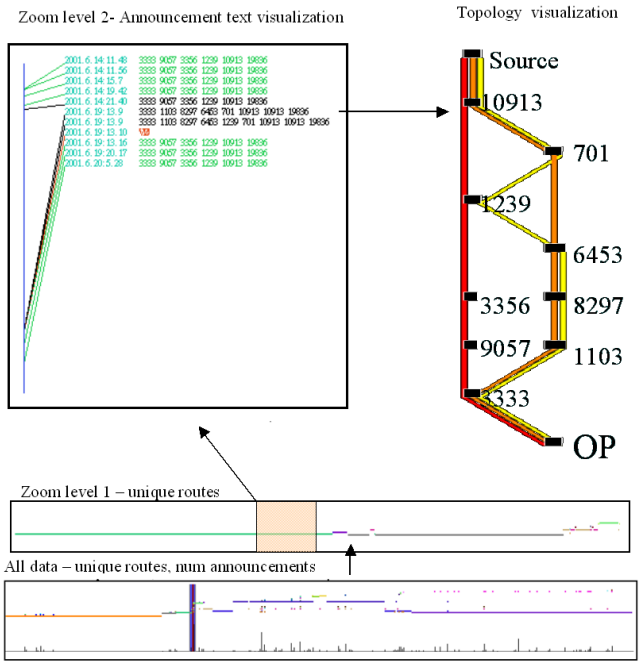

Figure 1: Layout of the different modules. The bottom row displays aggregate information of the entire time period (typically one year), and the user specifies a period to focus on in Zoom level 1. Unique routes visualization is used in Zoom level 1, where the user can specify another period of focus for Zoom level 2, which uses the text visualization and node/link visualization modules to display details of the BGP path announcements.

This visualization module is also used as a slider bar for the user to select periods of interest to focus deeper visual analysis. The user selects a window using mouse clicks, and slides the window over time with the mouse. In the detailed view, the user can further select a focus window to perform visual analysis using techniques described in following sections. This two-level overview+detail system allows an overview time period of one year, and detail period of several minutes, which is a very wide range of granularity. The layout of the two-level overview+detail interaction between the different modules is shown in Figure 1.

The plot of the number of announcements by itself can serve a very useful purpose. Figure 2 shows five spikes in the per-period count of the number of announcements. The user's attention is immediately attracted to these three events, and quickly focuses the detailed visualization tools on these periods. After investigating these events, the user then focuses attention on the other periods. This visual highlight of interesting events is in contrast to setting thresholds in non-visual analysis algorithms. Setting thresholds often results in sub-optimal identification of interesting events. It is challenging to select parameters such as bin size and which statistical measure to use. To overcome this difficulty, visualization-aided human judgment is used to identify reasonable cut-offs.

3.2 Route Announcement Module

Figure 3 shows BGP AS path announcements for one prefix from one peer as observed at a router. It is very tedious for the analyst to read the data and make sense out of it. Nevertheless, analyzing sequences of AS path announcements can potentially yield very valuable information. Therefore, our system provides browsing capability so that users can view these announcements in an intelligible

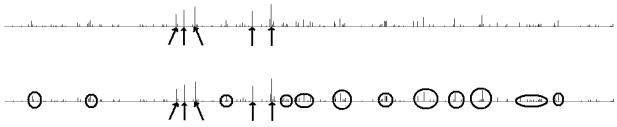

Figure 2: Visualization of the plot of the number of announcements in each unit period of time clearly shows five (arrowed) periods with extraordinarily large number of announcements. Other periods with more than usual the number of announcements are circled. It is more difficult for algorithmic (ie. automated, non-visual) programs to set optimal thresholds to flag periods of instability.

manner. Various text visualization methods and systems have been proposed in the past, such as the Information Visualizer [Card et al. April 1991], SeeSoft [Eick 1994] and the Document Lens [Robertson and Mackinlay November 1993]. However, the timestamp of each BGP announcement is of special importance, and none of these existing methods are designed to display such data. Therefore, a text visualization method has to be custom-designed to convey the time associated with each line of text.

Using the two-level overview+detail feature introduced in Section 3.1, the user can select a time interval. This time interval will be shown as a vertical line, called the *timeline*. Each AS path announcement within this time interval will be displayed chronologically. A line is drawn from the announcement to the time-line, indicating the timestamp of the announcement. In this way, clusters of announcements becomes obvious. We use the term *instability event* to describe a cluster of announcements, because the short-lived paths are unstable. Since the announcements occur in a cluster, they are probably due to the same cause and can be considered a single event. Our system is designed to pay particular attention to the discovery and analysis of such instability events because they are an indication of sub-optimal routing behavior and may reveal more serious flaws in the Internet architecture.

Figures 4 through 7 show some examples. Each AS path announcement is written as the timestamp (year month day hour minute) followed by the AS path. As shown in these figures, each unique AS path is written in a different color. This effectively emphasizes the different patterns shown in the figures. Sometimes, the user is not particularly interested in the individual AS numbers in the AS paths, and the colors assigned to some paths may be very similar. To make a clearer distinction between the different paths used so that the patterns may be clearer, we allow the user to choose a pictorial representation of each path, which is to represent each unique path as a box with a unique x-coordinate. Color is used as a secondary visual property to differentiate between boxes with close x-coordinates. Since *path* is a categorical attribute, proximity in color or x-coordinates does not have any meaning.

One problem with listing announcements sequentially and using lines to point to the timeline is that sometimes the boundaries between clusters may not be clear. It is beneficial to separate announcements more clearly. This is done in an alternative visualization method shown in Figure 8. In this display, the position of each announcement is determined by its timestamp. An offset is applied to announcements which are so close together that the letters would overlap.

Slow convergence, *oscillations* and *repeats* are phenomena which have been found and described by previous analyses of BGP data, such as [Gao and Rexford June 2000; Griffin and Wilfong Mar 2000; Griffin and Wilfong Aug 1999] and [Pei et al. June 2002]. Figure 6 shows an example of slow convergence, which occurs when a topology change is slow to propagate through the backup paths to the Observation Point. In this example, the link to the Origin AS has been broken, so a Withdraw announcement is

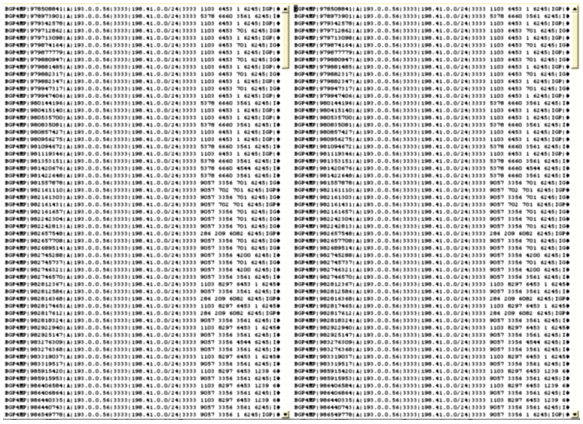

Figure 3: Analysis is difficult in simple textual display of BGP AS path announcements.

Figure 4: Oscillation is defined as closely-spaced consecutive announcements alternating between two unique paths.

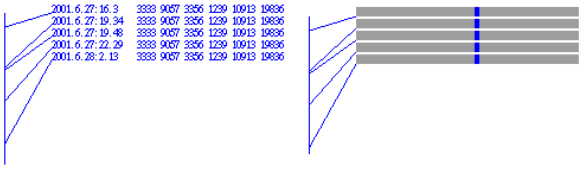

Figure 5: Repeats is defined as consecutive announcements of the same path.

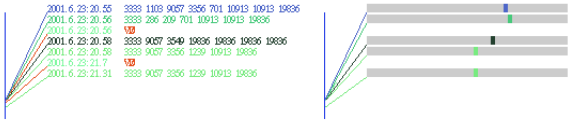

Figure 6: Example of Slow convergence. Slow convergence, defined as closely-spaced consecutive announcements of at least three distinct paths, occurs when a topology change is slow to propagate through the backup paths to the Observation Point.

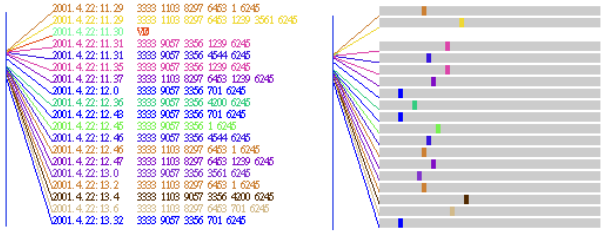

Figure 7: In this closely-spaced sequence of AS path announcements, features of both oscillations and slow convergence are observed.

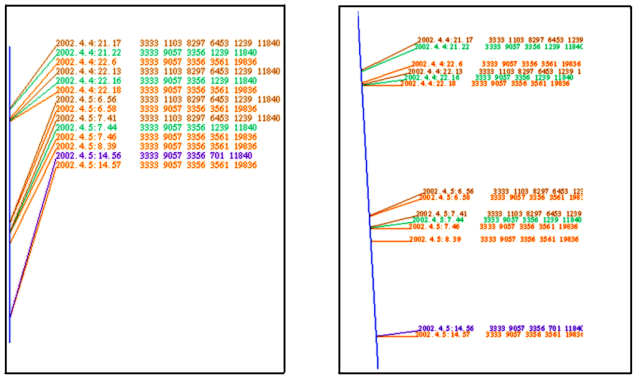

Figure 8: Two modes for the display AS path announcements. In display mode shown in the left picture is more orderly but the timestamps of the announcements are less obvious. In the mode shown in the right picture, each announcement is written at a position according to its timestamp. To prevent overlap between adjacent announcements, a slight offset is used to separate them.

sent through the primary path. However, the Withdraw announcement is slow to propagate through the routers in the backup paths, so AS-3333 announces the (already invalid) backup path (3333, 1103, 9057, 3356, 701, 10913, 10913, 19836) after the primary path has been withdrawn. When the Withdraw announcement has propagated through this path, AS-3333 then announces yet another backup path (3333, 286, 209, 701, 10913, 10913, 19836). When the Withdraw announcement has finally also propagated through this path, there are no valid paths left, so AS-3333 announces a Withdraw for this IP prefix. Our contribution is not the discovery of these phenomena but to ease the browsing of BGP data so that the user can efficiently examine lengthy sequences of BGP announcements and quickly observe instances of these events. Furthermore, some instances of instability events are worth further examination. The following sections describe some means of further analysis and the insights they yield.

Some new patterns have been found by interacting with the browsing feature, such as multiple-withdraw (shown in Figure 6), hybrids (shown in Figure 7), slow convergence after a withdraw, and sub-classifications of oscillations into “balanced” and “unbalanced”, and repeats into “short” and “long”. The identification and classifications of different events are important because they guide the focus of Internet analysts. Different patterns imply different causes, and more specific features of patterns further narrow the possible causes.

The browsing feature also conveniently addresses a problem encountered by signature-based event detection systems: What is the time-interval threshold to determine whether we consider a sequence of announcements part of the same “event”? Interactive visualization of the announcements allows users to see clustering patterns, as shown in some of the figures, and the timestamps of those announcements indicate what the threshold should be. One example of how human judgment is used effectively to find the threshold is to compare the cluster in Figure 5 and the one in Figure 7. The time intervals between the announcements in the cluster in Figure 5 are much longer than those in Figure 7, but browsing reveals clearly that both should be considered clusters. Visualization thus indicates that different thresholds should be applied for different event classes.

3.3 Node/Link Visualization Module

After instability events have been identified and classified, there is a need to analyze them further to find out which routers or links were responsible for the instability observed. In addition, we want to identify routers which may be chronically problematic. To perform such analysis, we need to build a visualization tool to observe the individual nodes and links in the path changes. Existing network visualization tools include NAM [Estrin et al. 1999], Otter [Huffaker et al. 1999], and the Internet Mapping Project [Cheswick et al. 2000], but there has been no technique that is custom-designed for detailed analysis individual nodes and links in path changes.

Designing appropriate visualization to convey the use of the nodes and links is a challenging task, so some of the methods we have experimented with were not very effective. For example, we found that animating AS path changes is not helpful in analysis because the user needs a persistent image of the different paths used. The most effective way we have found is to allow the user to click on AS paths in the text display to highlight them, and these paths are shown all together at the same time, so that the user can compare them. We designed two layout schemes to visualize changes in AS paths for the purpose of examining individual nodes and links.

The first method simply lays out the nodes and uses color to differentiate between the different paths used over time. The simple layout algorithm we employ arranges the nodes in the first path in a vertical line, and all new nodes in subsequent paths in vertical lines to the right. This is shown in Figure 9. From the figure, the user can see which nodes/links go in and out of use. This layout allows the user to conveniently analyze a small number of path changes. Unfortunately, it does not scale well; when more than four paths need to be visualized, the display becomes confusing. Therefore, we designed a second method, Arcs visualization, to analyze a larger number of path changes.

In Arcs visualization, each highlighted path is shown as a horizontal line of nodes joined by arcs. Each node represents one AS in the path, and is labeled by the AS number. The program assigns each AS the same unique x-coordinate for all paths in such a way as to avoid any backward (right to left) links. To differentiate between nodes with similar x-coordinates, color is used as a secondary visual mapping. The advantage of this visualization is that it is obvious when particular nodes are used, and when they start and stop being used. Drawing arcs instead of straight lines help makes it clearer when particular links are used. Figures 10 and 11 show how Arcs are used to assist in the challenging task of analyzing a sequence of path announcements. Figure 12 shows that perception about when particular links become used, unused and reused is less clear without arcs and colors.

3.4 Peer-Comparison Module

When instability events (such as those mentioned in Section 3.2) are observed, the network analyst would often like to conduct further investigation to find out the cause of the instability. Finding the cause helps to discover problems with the BGP architecture or problems with individual routers or links. One way to gather more clues about the causes of the observed problems is to compare the AS route path announcements at the same observation point for the same prefix from different peers.

There are three parts of peer-comparison. First, we compare the plots of the number of announcements of the same IP prefix from different peers, to determine if there are any similarities in the patterns. Next, we compare the announcements from the different peers. In both cases, similarities indicate that the announcements were caused by common nodes/links used by the paths from the different peers. The third part is to superimpose the paths on a topology graph to visualize the common nodes/links used. Examples of

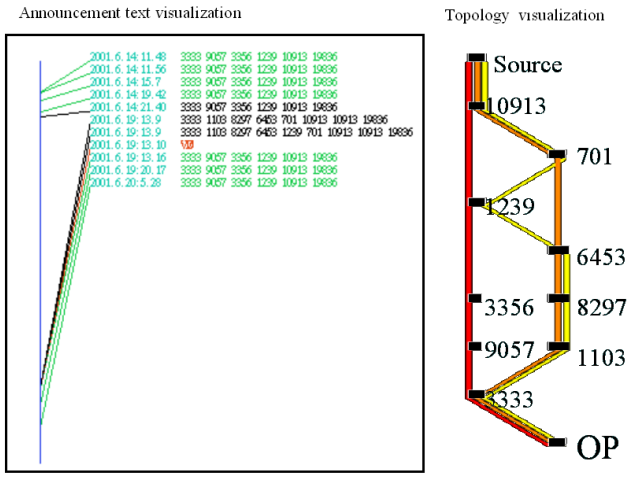

Figure 9: 2D Topology Visualization. Left: The three paths leading to the Withdraw are highlighted. Right: The first path is shown in red, the next in orange, and the last in yellow. From this sequence, the user infers that AS-10913 first sends a withdraw message to its peers, and the message arrives first through AS-9057. AS-3333 therefore switches to its backup path through AS-1103. The orange path is then withdrawn, and finally, the yellow path is withdrawn. The user infers that the slowness of the orange and yellow paths is due to the path between AS-6453 and AS-3333.

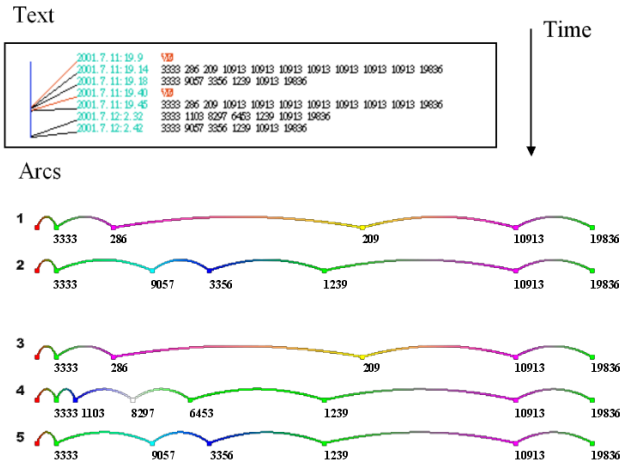

Figure 10: Path Visualization with Arcs to examine individual nodes and links. Each highlighted path is shown as a horizontal line of nodes joined by arcs. Each node represents one AS in the path, and is labeled by the AS number. Each AS is assigned the same unique x-coordinate for all paths. To differentiate between nodes with similar x-coordinates, color is used as a secondary visual mapping. The user has selected to highlight the five paths in the text display. This example shows two post-withdraw slow convergence cases. In such cases, the paths used earlier are considered "faster", because their announcement arrived first, whereas the paths used later are considered "preferred" because when their announcements arrived, they replaced the earlier paths. In this example, therefore, Path 1 is faster than Path 2, and Path 2 is preferred to Path 1. This conclusion is consistent with the second observed case (Paths 3 through 5) since Path 3 is identical with Path 1, and Path 5 is identical with Path 2. However, in the second case, a new path, Path 4, is slightly faster than Path 5 but Path 5 is preferred. Paths 4 and 5 have node AS-1239 in common. These inferences are further used to interpret the observations shown in Figure 11.

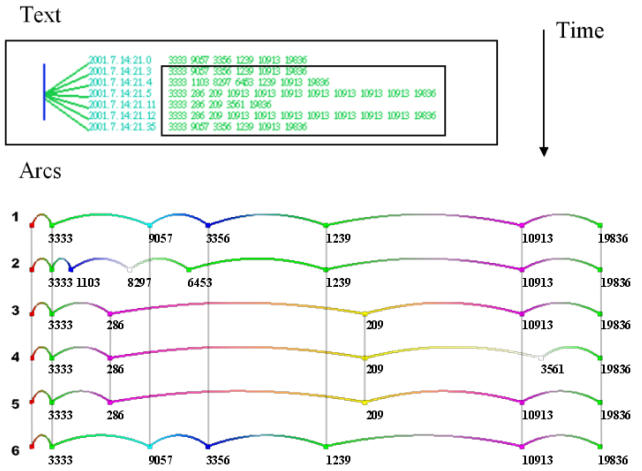

Figure 11: Interpretation of a complex sequence using Arcs. The user has selected to highlight the last six paths in the text display. Assuming that one single point of failure caused this sequence of announcements, the user infers that the point of failure is either the source AS (AS-19336) or the link 19336-10913. The change from Path 1 to Path 2 is caused be the slow arrival of the "withdraw" message from Path 2. This is consistent with the inference from Figure 10 that Path 1 is preferred, but Path 2 is sometimes faster and sometimes slower. By the time the Path 2 is withdrawn, AS-10913 has announced the prefix again, causing the Path 3 to be used. This assumes that Path 3 is faster than Paths 1 and 2, which is consistent with Figure 10. Some time before 21:11, the prefix is once again withdrawn, causing the backup Path 4 to be used. The prefix is then announced again through AS-10913, leading to Path 5. Path 6, which is the preferred path, is late in arriving, which is consistent with the observations in Figure 10. This hypothesis asserts that the prefix was withdrawn twice in this short period of time, even though the "withdraw" message did not propagate to the observation point. The hypothesis is plausible because just days prior, two withdraws also happened within a short period if time, as shown in Figure 10. The challenging task of explaining an unusual pattern of path announcements is thus accomplished with the help of Arcs visualization and domain knowledge. Additionally, the user has selected to display vertical lines to help the user visually align the points on different paths representing the same node.

Figure 12: Straight line representation of paths in Figure 11. Use and re-use of individual links are less obvious in this representation. This shows that colors and arcs are helpful in accentuating the perception of links.

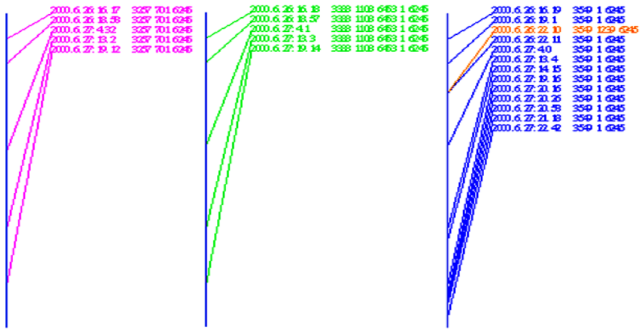

Figure 13: Comparison of BGP AS path announcements from three peers, AS-3257, 3333, and 3549, to find the reasons for the announcements. The similarity in the timing pattern of the announcements from all three peers suggests that the origin of the repeated announcements of the same path is the origin AS itself.

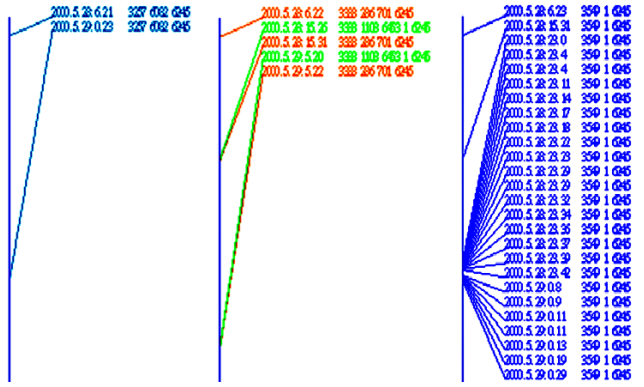

Figure 14: The presence of the large number of repeated announcements of the same route from only Peer 3549 suggests that the source of the problem is related to AS-3549. The additional announcements from AS-3333 not present in AS-3257 is probably due to propagation delay because of the short interval between the green and the orange paths.

insights gained from such a comparison are shown in Figures 13 through 15.

3.5 Events Visualization Module

In Section 3.2, we have shown some examples of different *types* of instability events, such as slow convergence and oscillations. The classification of events into types is based on the user's domain knowledge and insights from visual analysis. It would be useful for the user to have a summary of the instability events. In particular, the user would like to know the distribution of the various types of events, how many events are hybrids, which events occur more frequently and which events are more severe. In the Events Visualization module, the goal is to present a visual summary of the instability events in meaningful and insightful manner.

We use *EventShrubs*, a new visual metaphor we designed specifically for presenting *instability events*. An event is defined as a 4-tuple (*starttime*, *endtime*, *size*, *flags*). For visualizing instability events in BGP AS path announcements, the *starttime* and *endtime* are simply the timestamps of the first and last announcement respectively in the instability event. The *size* parameter is the number of announcements in the instability event, and each *flag* corresponds to one class of instability event, such as "repeats". A flag

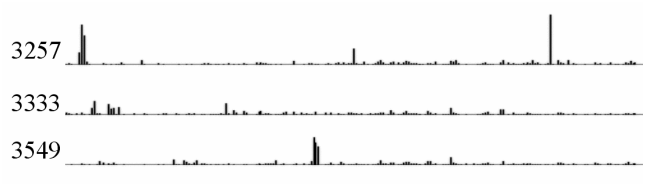

Figure 15: Comparison of the count of the number of announcements from each peer shows no correlation. This is because the "spikes" are dominated by the large instability events. Correlated patterns during stable periods such as those shown in Figure 13 are obscured in this view. The user infers that instability events are caused by intermediate nodes in the AS paths.

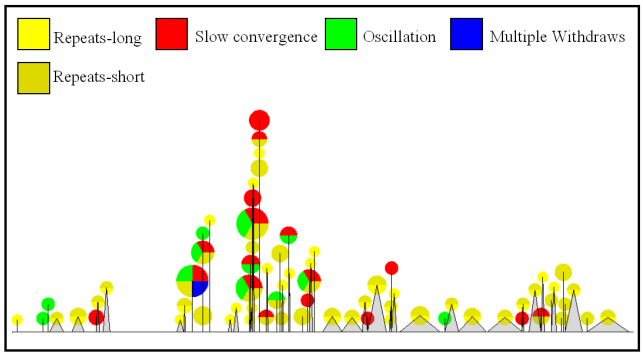

Figure 16: EventShrubs used to visualize instability events. Each "shrub" consists of a crown (circle) and a trunk (vertical line or triangle). Each segment of a circle indicates a type of instability event, colored according to the legend in the upper part of the picture. The presence of a colored segment in a circle represents the presence of characteristics of the type in the instability event. The area of the circle represents the number of announcements in the event, giving more emphasis to the more severe events. The base of each shrub covers the duration of the event. The height of each shrub has no meaning, and is used only to prevent occlusion. Nevertheless, tall shrubs indicate that there are many events at that time. EventShrubs give users a clear presentation of the distribution of different events, their severity, duration, and frequency, and also effectively shows hybrids of different types.

is turned on for an event if the event shows characteristics of that class. Figure 16 shows and explains the visual metaphor used in EventShrubs to visualize instability events. Figure 17 compares some EventShrubs generated for AS path announcements of different IP prefixes from two different peers. The obvious differences in the characteristics of the different EventShrubs reveal problems with specific ASes and/or links and direct the network analyst to focus further investigation.

From these figures and their accompanying explanations, the user can verify that the EventShrubs metaphor is effective in conveying a rich set of information regarding the distribution of the different types of events. These pictures help network analysts better understand the occurrences of the instability events.

4 Conclusions

We have used visualization for the analysis of BGP AS path announcements, which lies at the foundation of Internet connectivity. Our goal in analyzing BGP announcements is to learn about the vulnerabilities of the Internet architecture to improve it security. The

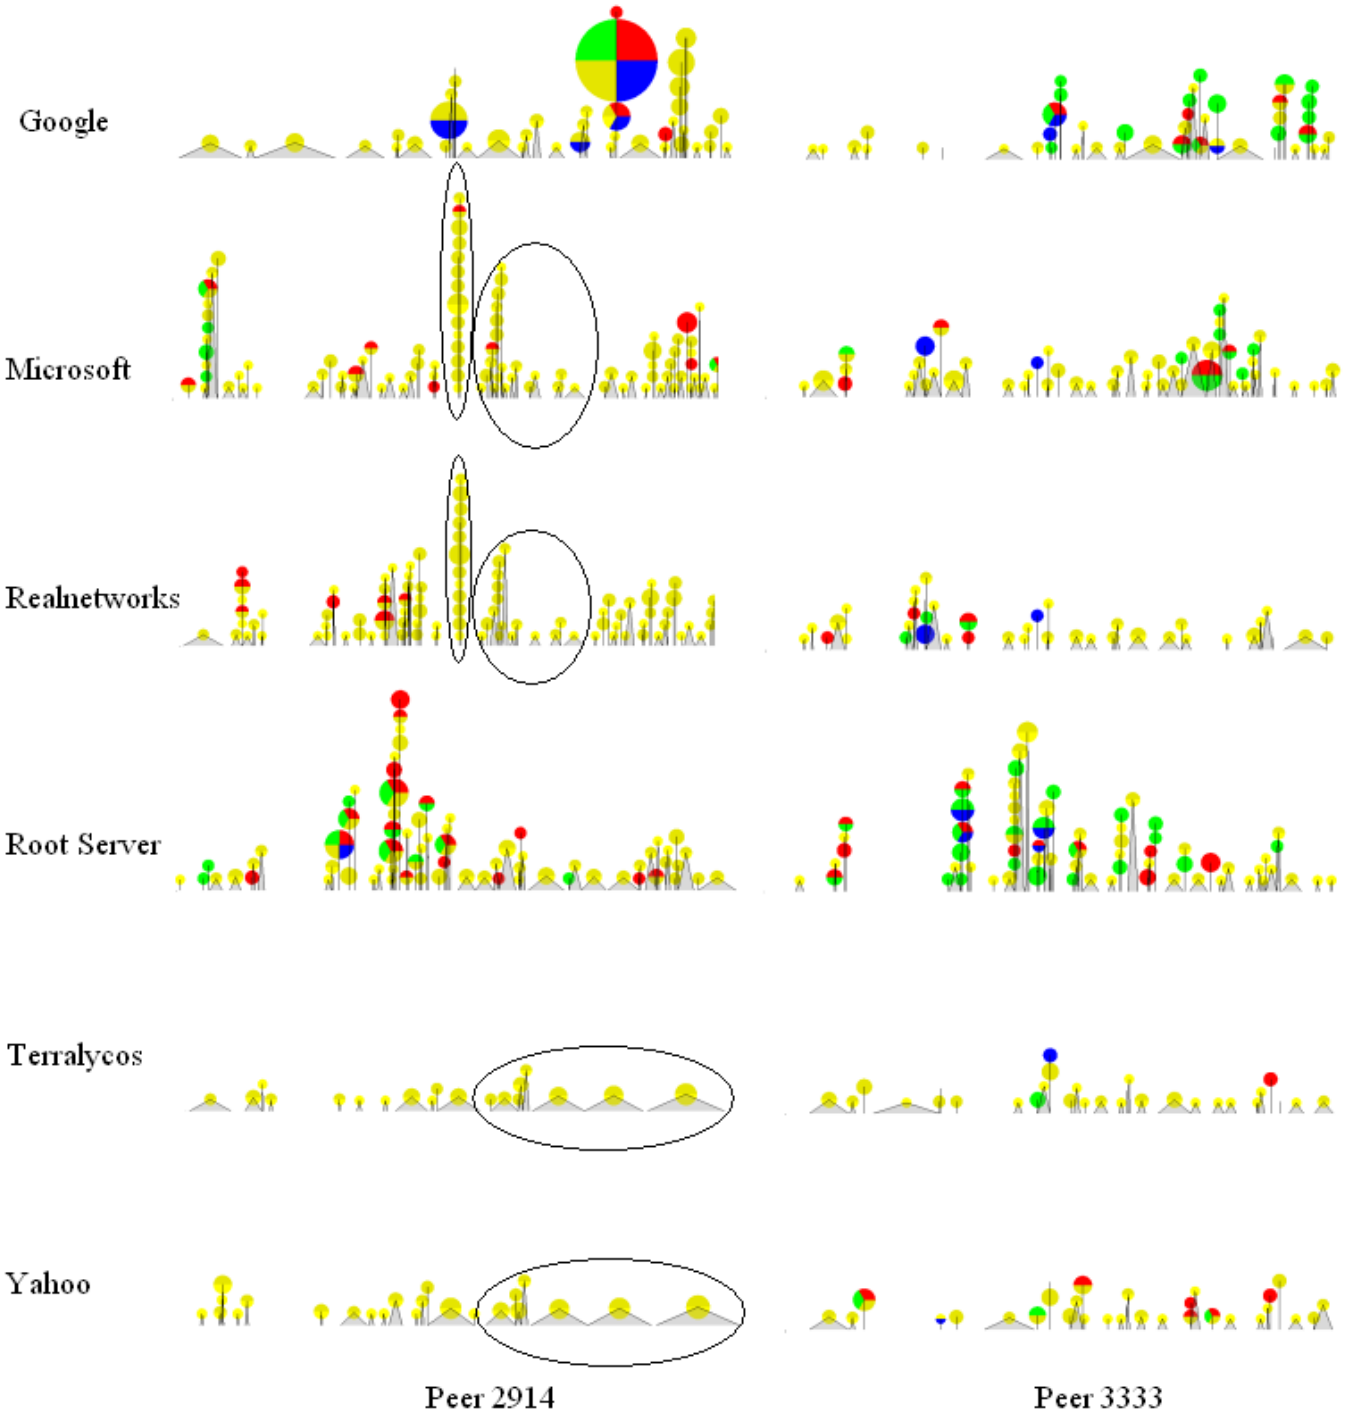

Figure 17: Comparing the different EventShrubs generated from five different IP prefixes and from two different peers. The EventShrubs of Google and Root Server especially stand out because of the number and severity of hybrid instability events. Similarities are also worth attention. For example, Terralycos and Yahoo from peer AS 2914 show similarity, and Microsoft and Realnetworks from AS 2914 also show similarity in circled areas. Triggered by these observations, further investigation (not described in this paper) of these events reveal that they are indeed related. There is little observable correspondence between the instability events associated with the same IP prefix but from two different peers.

visualization demands of Internet routing analysis require many different modules, each targeting different aspects of the analysis.

We have presented the five modules we use in our system, and have outlined a sequence for their use, from general views to very specific and detailed analyses, so that the user can gain a comprehensive picture of routing dynamics.

From our experience in using the system, the visual analysis of BGP contributes to the security and stability of the Internet in two main ways: *inquiry* and *diagnosis*.

Inquiry refers to the examination of BGP data to look for any signs of problems in Internet connectivity. We have shown various examples of how the aggregate visualization (number of announcements plot and unique paths) module, the announcement text visualization module and the node/link visualization module work together to discover and identify instability events. From interactive visualization, the user has been able to detect instances of known types of instability events as well as identify new types, and further classify known types. That is a major contribution to the understanding of BGP dynamics.

Diagnosis refers to the deeper investigation and analysis of the nature and causes of a problem once it has been discovered. We have shown how the highlighting of AS paths in the topology display, the comparison of announcements from different peers and the visualization of the summary of instability events with EventShrubs are able to reveal much information such as the source of the observed problems. We have described the insights and knowledge gained by the user, allowing inferences of the causes of the instability events. For example, the user was able to quickly identify faster paths, preferred paths, slow points, points of failure, as well as hypothesize about the sequence of events that led to the observations. The user also discovered new sub-classes of instability events through visualization. All these insights and discoveries have been sought after by routing analysts because they help develop understanding of Internet routing dynamics. It is not impossible for the tasks to have been performed without the help of visualization, but visualization made it much easier for the user to obtain these important insights.

We have received very positive feedback from the network analysts and Internet researchers who have seen and used the visualization software. In the continuation of our work, we would like to release the software to network operators. More evaluation of the system by operators and analysts would guide the further development of the visualization system. As more analysis is being conducted on the data, we expect more insights to be revealed and more discovery of Internet vulnerabilities. Extensions to the visualization software then need to be made to focus on the analysis of these new discoveries.

Besides the contributions to Internet security, this applications paper also contains several ideas relevant to information visualization. First, we have effectively integrated a two-level overview+detail system with different visualization modules. Second, we have presented a simple text visualization technique that is customized for visualizing text with timestamps. This is an example of how colors and positions can effectively reveal patterns. Third, we have designed Arcs visualization to help perceive path changes. Fourth, we have introduced EventShrubs, a new metaphor for visualizing events. We have shown how the simple and compact figures in EventShrubs are able to convey rich information. Although the visualization techniques introduced in our system were designed for Internet routing analysis, they are likely to be usable for other applications.

Interactive visualization tools have not been widely used in Internet security efforts. We have built a foundation for greater incorporation and utilization of human visualization and interaction for learning about the Internet architecture to make it more stable and secure. It is our hope that the insights drawn from visual analysis would eventually lead to recommendations for improving the protocols and administrative policies currently used in the Internet.

5 Acknowledgments

This work has been sponsored in part by the U.S. National Science Foundation under contracts ACI 9983641 (PECASE award), ACI 0222991, and ACI 0220147 (ITR). The authors would also like to thank Ke Zhang for preparing the test data set, and TJ Jankun-Kelly and Brett Wilson for proofreading an early version of this paper.

References

CAIDA. The Cooperative Association for Internet Data Analysis. *http://www.caida.org*.

CARD, S., ROBERTSON, G., AND MACKINLAY, D. April 1991. The Information Visualizer, and Information Workspace. In *Proc. CHI '91*, 181–188.

CHESWICK, B., BURCH, H., AND BRANIGAN, S. 2000. Mapping and Visualizing the Internet. In *Proc. 2000 USENIX Annual Technical Conference*.

EICK, S. 1994. Graphically Displaying Text. *Journal of Computational and Graphical Statistics 3*, 2, 127–142.

ESTRIN, D., HANDLEY, M., HEIDEMANN, J., MCCANNE, S., XU, Y., AND YU, H. 1999. Network Visualization with NAM, the VINT Network Administrator. *Tech. Report 99-703, Univ. of Southern California*.

GAO, L., AND REXFORD, J. June 2000. Stable Internet Routing Without Global Coordination. In *Proc. ACM SIGMETRICS*, 307–317.

GIRARDIN, L. 1999. An Eye on Network Intruder-Administrator Shootouts. USENIX Assoc, Berkeley, CA, USA, 19–28.

GRIFFIN, T., AND WILFONG, G. Aug 1999. An Analysis of BGP Convergence Properties. *Proc. ACM SIGCOMM*, 277–288.

GRIFFIN, T., AND WILFONG, G. Mar 2000. A Safe Path Vector Protocol. *Proc. IEEE INFOCOM*, 490–499.

HUFFAKER, B., NEMETH, E., AND CLAFFY, K. 1999. Otter: A General-Purpose Network Visualization Tool. In *Proc. INET '99*.

IETF. The Internet Engineering Task Force. *http://www.ietf.org*.

MYER, D. University of Oregon Route Views Project. *http://www.antc.uoregon.edu/route-views/*.

PEI, D., ZHAO, X., WANG, L., MASSEY, D., MANKIN, A., WU, S., AND ZHANG, L. June 2002. Improving BGP Convergence through Consistency Assertions. In *Proc. IEEE Infocom*.

REKHTER, Y., AND LI, T. 1995. A Border Gateway Protocol 4 (BGP-4). *RFC 1771*.

ROBERTSON, G., AND MACKINLAY, D. November 1993. The Document Lens. In *Proc. UIST '93*, 101–108.

TEOH, S., MA, K.-L., WU, S., AND ZHAO, X. 2002. Case Study: Interactive Visualization for Internet Security. In *Proc. IEEE Visualization*, 505–508.