

LOCAL-AREA PATH DIVERSITY IN THE INTERNET

Wei-haw Chuang, Greg Johnson, Aditya Ojha, Nadya Williams and Joy Xin
University of California San Diego
9500 Gilman Drive, La Jolla, California 92093
email: wchuang, aojha, nadya, qxin@cs.ucsd.edu, johnson@sdsc.edu

Abstract

The Internet is a rapidly growing, dynamic system. Several thousand new hosts are added daily [1]. Worse, streaming media applications - often requiring better than best-effort quality of service - are becoming increasingly popular. For this growth to be sustainable, such change must be met with an equally explosive increase in capacity, and improved utilization of existing networking infrastructure. In this report, we examine the latter, characterizing the presence of path diversity in a randomly sampled collection of hosts, and presenting evidence that this diversity is (to a limited degree) being exploited to distribute traffic load.

1 INTRODUCTION

The population of the Internet is on the order of several hundred million hosts and counting. In addition, the demand for applications which require a level of quality of service beyond best-effort, such as streaming audio / video content, is on the rise. However, the installation of new networking infrastructure to meet this increasing demand, is complicated by issues of cost and scalability. Viewing the Internet as a hierarchical structure of backbones, domains, networks, and subnetworks, the addition of capacity at one level suggests that capacity at higher levels must be increased by a similar amount. These problems point to finding new techniques for improving the utilization of existing resources. As a result, congestion and the ability to mitigate its effects by spreading traffic over several routes to the same destination is an area of intense interest.

In this text, we report on the topographic character of path diversity in a randomly sampled set of hosts, based on data obtained from the Cooperative Association for Internet Data Analysis (CAIDA). We then present an analysis of this data, showing the presence, frequency of occurrence, temporal behavior, and effectiveness of load balancing in multi-path routes.

2 A LOAD BALANCING PRIMER

Load balancing is the concept of distributing service requests over multiple physical resources to minimize contention for any single functional unit. In the context of computer networks, packet traffic might be distributed across several unique paths to the same destination in an effort to avoid congestion on any single link. There are several reasons to study load balancing. For example, single-path routing algorithms like the popular shortest-path routing protocols have the disadvantage of being unable to switch routes when the "shortest" becomes congested. Thus, packets will be lost, even while other paths to the same destination go unloaded. Also, multimedia applications like video conferencing, remote medical imaging, and video-on-demand, have stringent delay, throughput and packet-loss rate requirements.

Load balancing allows the network to take advantage of multiple best paths to a given destination (and end-host or an intermediate router). Packets (possibly from the same application) are sent over several paths so as to attain an even spread of traffic through available links and nodes. Doing so can mitigate the effects of congestion, by taking advantage of capacity on links that would otherwise be idle. Load balancing can be conducted on a per-destination or per-packet basis. With per-destination load balancing, the router distributes packets based on the destination address. Given two paths through the same region of the network, all packets for destination 1 are routed over the first path, all packets for destination 2 are routed over the second path. Alternatively, with per-packet load balancing the router spreads the packets for any given destination across all available paths.

Several multi-path routing techniques have been proposed. These include "multiple disjoint paths" [14] based on distance-vector routing, "dynamic multi-path routing" [3] based on source routing, and randomized multi-path routing schemes [4, 5]. Additionally, there are several approaches for determining how traffic should be distributed across multiple paths. These load balancing algorithms may be either static or dynamic. Static




Figure 1: A simple example of a "diamond" topology with three different paths to the same node.

schemes are those which assign an unchanging quantity of traffic to each link, and include round-robin and randomization [9]. Unlike static schemes dynamic algorithms are able to adapt to the changing availability of capacity. For example, a node might route the next packet across the least loaded link. Other examples include "least connections" and "fastest response".

3 LOCAL-AREA PATH DIVERSITY

In this report we are primarily concerned with detecting and characterizing path diversity (and any associated load balancing), in a simple network topology hereafter referred to as a "diamond". A diamond is represented by two nodes connected via several dual-hop paths. That is, given nodes A and Z , two or more paths fan out from A to intermediate nodes (siblings) C, D, \dots , and paths from these nodes fan back into Z . An example of this topology with three paths to the same node is shown in Figure 1. We have verified the existence of diamonds containing up to twenty siblings.

Note that we refer to path diversity and load balancing in this topology to be "local-area", as path traversal and load balancing decisions are made in the context of this local region only. No regard is given to the end-to-end route traversed by any given traffic. Measuring the presence (or lack) of load balancing in this topology is a matter of detecting how often traffic is routed over each of the paths belonging to a given diamond. The presence of a diamond in the network does not of course imply that all paths in that diamond are exploited equally.

4 SKITTER DATA

The dataset from which our analysis is derived comes from the Cooperative Association for Internet Data Analysis (CAIDA), and was collected using a tool called Skitter. Skitter actively probes hosts on the Internet and gathers ICMP TTL data, which may then be used for mapping and analyzing network topology and performance

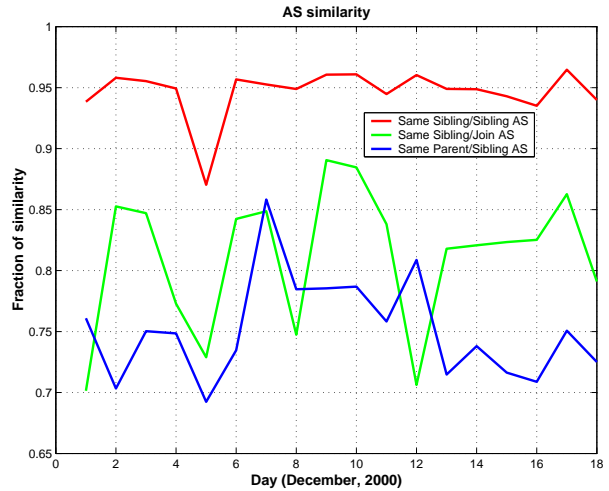


Figure 2: A graph showing the AS domain relationship between sibling, parent (fan-out source), and join (fan-in destination) nodes.

[10]. Skitter tracks forward IP paths (the "hops") from a source to many destinations in a manner similar to that of traceroute. It collects round trip times (RTT) and traces paths by incrementally increasing the TTL of subsequent probe packets and recording the hops indicated by the corresponding ICMP responses.

The data used here consists of probe results for 1793 geographically diverse hosts. Each destination was probed at the rate of six times per hour from the Urbana-Champaign US monitor (champagne.caida.org, 141.142.121.4), between December 1 and December 18, 2000. These raw results were preprocessed by a CAIDA tool, producing a file containing source and destination addresses, probe timestamps, RTTs, and the number of hops to a given destination. Our team further processed the data (with scripts customized from CAIDA originals) to reveal the existence and degree of path diversity and load balancing. We first looked for the presence of diamonds and the frequency of occurrence for diamonds with two or more paths. Next we attempted to ascertain the presence and character of load balancing by examining the distribution of probes along each path in a given diamond, and how these paths shifted over time. Further, we developed a script to convert Skitter trace data into graphs showing the topology of the network between a given source and destination. The script relates the trace data to geospatial reference and parent AS (autonomous system) data also provided by CAIDA. In doing so, we uncovered two important characteristics of the placement of siblings belonging to specific diamonds, relative to AS domain and country boundaries.

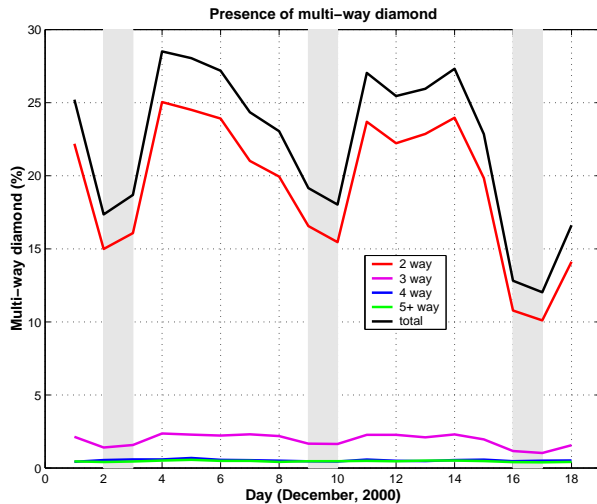


Figure 3: Percentage of diamonds with two or more paths, measured during December 1-18, 2000. Shaded areas correspond to weekends.

5 ANALYSIS

Figure 9 illustrates two key points regarding the relationship between diamonds and the physical location of their sibling hosts. Each oval represents a node (labeled with its IP address and associated DNS name). Source and destination nodes are identified by a double circle, and boxes enclose all nodes belonging to a given AS. Graph edges are colored based on the number of times each was traversed by a probe. Black edges have probe counts greater than 100, blue edges less than 100, and green edges denote situations in which Skitter probes could not reach the destination.

Two subgraph structures stand out. The first comprises the VERIO and TELEGLOBE AS regions. Of the four paths between the two domains, a single route dominates. The remaining three paths may be due to routing table flutter, possibly as the result of BGP updates as suggested in [11]. The second structure is composed of the APNIC-AS-3-BLOCK AS cluster, and an unnamed “-” AS. Between these we find an example of a six-way diamond across which traffic is relatively balanced, suggesting the presence of a load balancing scheme [9]. An interesting property of this diamond is that both its source (203.160.227.252) and siblings lie within the same domain, while the destination (10.6.1.1) is located in another. Figure 2 shows that this case is not unique. Specifically, sibling nodes belong to the same AS domain nearly 95% parent (fan-out source) and join (fan-in destination) nodes are located in the same domain as siblings

75

Thus far we have demonstrated that path diversity does

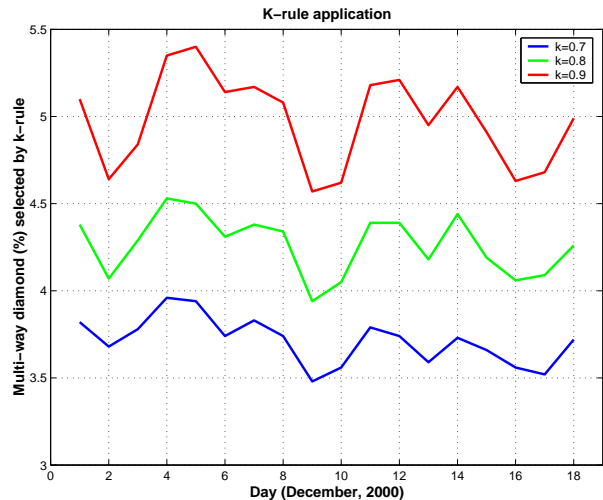


Figure 4: Selection of diamonds by the k -rule from the total.

exist, and that it may potentially be employed for load balancing. In the next following subsections we examine the nature of path diversity and load balancing in detail.

5.1 Path Diversity

We analyzed 383734 diamonds for 1793 destinations. For each diamond we counted the number of different paths by which its destination node could be reached. We found that on average 22% of all diamonds have more than one path. Among these multi-path diamonds 2-path diamonds represent the majority, or 86%, 3-path diamonds constitute only about 9%, with the remainder shared mostly among diamonds with 4, 5, or 6 paths. Further, we found a very small number of diamonds with a surprisingly high number of paths: on average 10 diamonds with 8 paths, 2 with 9 paths, two with 10, and 1 (each) with 15, 16, and 20 paths. Figure 3 shows this distribution over the course of the 18 day period. Shaded bars correspond to weekends.

Note that the presence of diamonds increases during the high-traffic weekday periods, and that this cyclic change is most evident in 2 and 3-path diamonds. It is difficult to assess the reliability of this correlation for diamonds with higher (4 or more) path counts, since the overall fraction of diamonds in this category is small. It may be necessary to analyze such diamonds over a longer duration to see a distinct pattern. Nevertheless, the persistence of such diamonds shows that they are a stable feature rather than a fluctuation or unusual behavior.

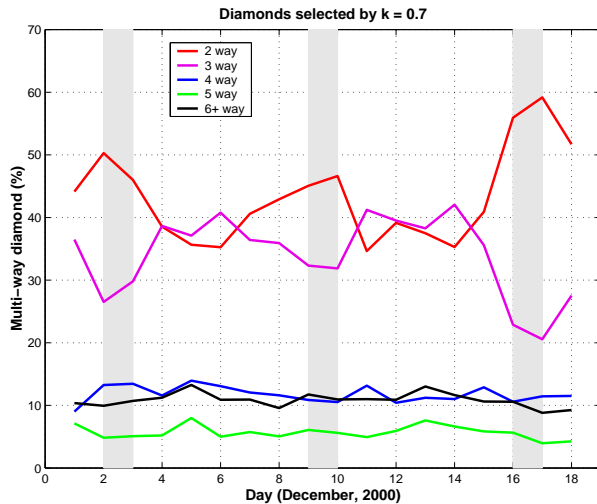


Figure 5: Percentage of diamonds showing "good" load balancing ($k = 0.7$) during December 1-18, 2000. Shaded areas correspond to weekends.

5.2 Load Balancing

Having identified the extent to which path diversity exists in the data, we now turn our attention to characterizing the subset of diamonds which appear to be utilized for load balancing. Load balancing in this context involves dividing traffic among multiple paths to alleviate congestion. To find out how balanced the paths in a given diamonds are, we applied the following rule: no more than a k th-fraction of the total probe traffic through a diamond should pass through any single path. For this calculation we selected three different values for k : 0.7, 0.8 and 0.9.

These results are shown in Figure 4, and demonstrate that "good" load balancing (no single link carries more than 70% of a diamond) is not all that common. On average, diamonds in this category constitute only about 3.7% of the total diamond count. In other words, for the majority of diamonds, it is more common for a few paths to handle the bulk of the traffic flow.

Moreover, the majority of diamonds in which we consider the load to be well balanced, most of these diamonds contain only two or three paths. Figure 5 summarizes these results. Interestingly, the two and three-path diamonds show daily fluctuations that are nearly "mirror" images of one another. At this time we are unable to explain this phenomenon.

To find out what load balancing achieves once implemented, we examined three multi-path diamonds in detail (one each with 3, 5, and 6 paths). These results are shown in figures 6, 7, and 8 respectively.

In the 3-path diamond, one of the paths is used twice as

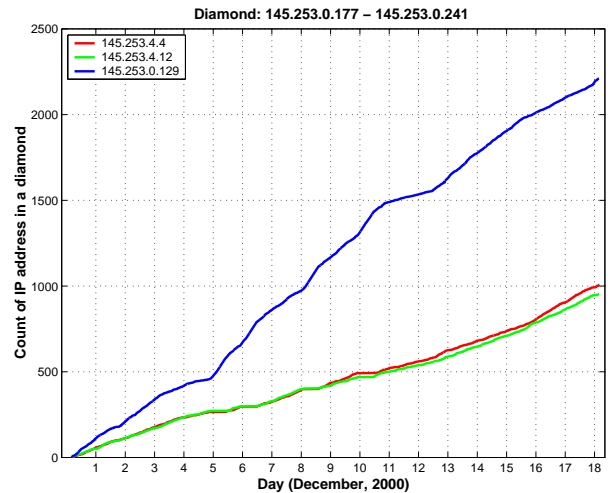


Figure 6: Accumulation of traffic over each path in a 3-way diamond during December 1-18, 2000.

often as either of the remaining two. The load across these two links appears well balanced, while the load of the third looks to be consistently twice the others, and may be the result of dynamic load balancing. The outstanding path may result from either a higher bandwidth link, or from routing policy.

In the 5-path diamond, we observe a weighted load distribution. Two of the five paths show the lowest load (magenta and black lines overlap on the figure) and exhibit even distribution. Each of the remaining three has a next higher fraction of the total load. It appears as if the distribution is almost fixed by weighted coefficients, and the proportion of traffic allocated to each path in the diamond is constant.

In the 6-path diamond, the traffic load is consistently balanced, and is likely achieved by dynamic routing.

6 CONCLUSION

Local path diversity clearly exists and is consistently present, offering ample opportunity for load balancing. We note, however, that efficient load balancing is rarely employed, and consists of only about 3.7% of the cases where multiple paths exist to the same destination (end host or intermediate node). We've found evidence of the use of several load-balancing approaches including what appears to be weighted round-robin and dynamic balancing. We've found correlation between the use of load balancing and weekend periodicity. The number of 2-path diamonds used in the routes has a tendency to rise, while the number of 3-path diamonds has a tendency to

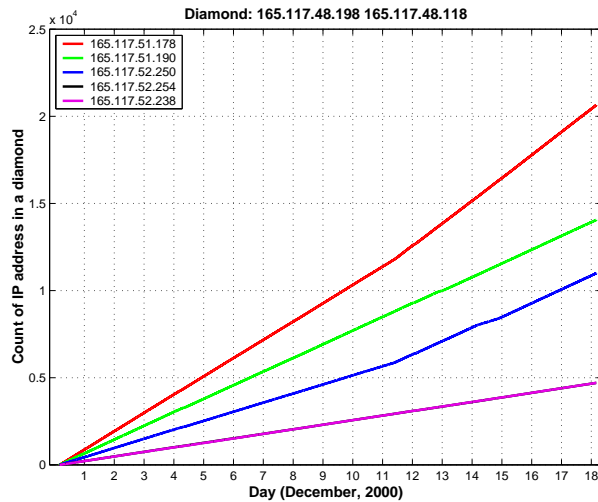


Figure 7: Accumulation of traffic over each path in a 5-way diamond during December 1-18, 2000.

fall during the weekend compare to the weekdays. Furthermore, the fall and the rise are at the expense of each other. We hypothesize that during the weekends, when the load is usually lower, there is less possibility for congestion and routing instability, and using 2 paths in place of 3 is sufficient. For the multi-path diamonds there is no such evidence. Approximately 95 which offers a good opportunity for load balancing since they are affected by a single AS routing policy but this balancing is by and large not yet implemented. We've found that non-US networks show a bias towards higher fan-out structures compare to that of US networks.

Since load-balancing distributes processing and communications activity evenly across network, it is especially an important technique to improve performance of the network. It increases the effective bandwidth between the source and the destination so that more data packets can be delivered. Load-balanced routing protocols tries to distribute traffic evenly over the whole network, and in consequence, increases network performance with respect to throughput, message loss, message delay and link utilization.

7 FUTURE WORK

Any collection of raw data is likely to have outliers due either to glitches in the programs used for data collection, or to problems with the equipment or application method. However rare these cases might be, it is always better to eliminate such outliers or inconsistencies before using the data for any further processing. When assessing the route stability for example, it would be useful to

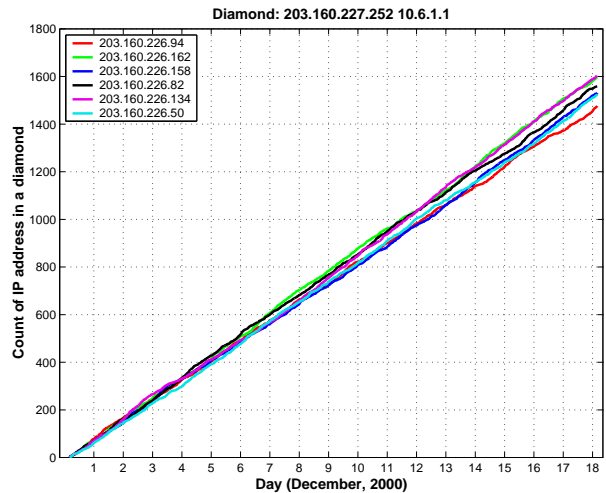


Figure 8: Accumulation of traffic over each path in a 6-way diamond during December 1-18, 2000.

determine if the router always returns the same address. A single router may have two different IP addresses but a single DNS entry. Alternatively, IP addresses can resolve to very similar names which, in fact may be the same router. For such situations it would be useful to "merge" these cases and use a single equivalent IP or DNS entry.

Skitter measurements used for our analysis are made on average every 10 minutes for each destination. We would not be able to use this approach to determine the existence of short-lived routes, because they may come to life and go dormant during this interval without ever being recorded. However, if such short-lived routes show up as alternative paths during a limited number of probes, they could be analyzed separately so as not to skew the distribution of diamonds. We cannot currently estimate to what degree this may affect our analysis. In order to find the presence of short-lived routes it would be necessary to collect Skitter measurements at a much higher frequency for "suspect" routes.

Both path diversity and the stability of load balancing depend on the time scale. Can we catch diurnal patterns, busy and off hours, and find out how similar or different they are for the geographically different regions? Do they follow the same pattern? It may be useful to separate the data into groups according to likely congestion during peak hours or different week days, and explore time-of-day or day-of-week variations in the presence of the diamonds, and the performance of the routes that use these diamonds.

Analysis of load balancing structures other than diamonds would help to generalize our study. Another direction would be to look further for any correlation be-

tween the geographic location of diamonds.

We would like to complete our study of low edge count paths and their negative impact on path diversity. So far we have had inconsistent results in this area of our analysis.

Acknowledgements

We are grateful to the following individuals for providing our team with data, advice on project topics, and patient instruction on the finer points of path diversity and load balancing in today's Internet. Special thanks to Andre Broido, Bradley Huffaker, and David Moore (CAIDA), Stefan Savage and Geoff Voelker (UCSD).

References

- [1] U. S. Department of Commerce. The Emerging Digital Economy. <http://www.ecommerce.gov/emerging.htm>, April 1998.
- [2] Load Balancing for Server NICs. http://www.3com.com/technology/tech_net/white_papers/503006.html, 2001.
- [3] S. Bahk and M. Zarki. Dynamic Multi-path Routing and How it Compares With Other Dynamic Routing Algorithms for High Speed Wide Area Networks. In *Proceedings of the 1992 ACM SIGCOMM Conference*, October 1992.
- [4] S. Bak and J. Cobb. Randomized Distance-Vector Routing Protocol. In *Applied Computing (ACM symposium '99 proceedings)*, February 1999.
- [5] S. Bak, J. Cobb and E. Leiss. Load-Balanced Routing via Randomization. In *Proceeding of CLEI*, August 1999.
- [6] A. Broido. Load Balancing and Load Sharing. <http://www.caida.org/broido/ldb/ldb.html>, 2001.
- [7] B. Chandra, M. Dahlin, L. Gao and A. Nayate. End-to-end WAN Service Availability. To appear: *Third Usenix Symposium on Internet Technologies and Systems (USITS01)*, January 2001.
- [8] B. Fortz and M. Thorup. Internet Traffic Engineering by Optimizing OSPF Weights. In *Proceedings of IEEE INFOCOM'2000*, March 2000.
- [9] E. Gustafsson and G. Karlsson. A Literature Survey on Traffic Dispersion. In *IEEE Network*, vol. 11 no. 2, March / April 1997.
- [10] B. Huffaker, M. Fomenkov, D. Moore and K. Claffy. Macroscopic Analyses of the Infrastructure: Measurement and Visualization of Internet Connectivity and Performance. To appear: *Proceedings of PAM2001*, April 2001.
- [11] C. Labovitz, A. Ahuja and F. Jahanian. Experimental Study of Internet Stability and Wide-Area Network Failures. In *Proceedings of IEEE FTCS*, June 1999.
- [12] V. Paxson. End-to-End Routing Behavior in the Internet. In *Proceedings of the 1996 ACM SIGCOMM Conference*, August 1996.
- [13] S. Savage, A. Collins, E. Hoffman, J. Snell and T. Anderson. The End-to-end Effects of Internet Path Selection. In *Proceedings of the 1999 ACM SIGCOMM Conference*, September 1999.
- [14] D. Sidhu, R. Nair and S. Abdallah. Finding Disjoint Paths in Networks. In *Proceedings of the 1991 ACM SIGCOMM Conference*, September 1991.
- [15] H. Tangmunarunkit, R. Govindan, S. Shenker and D. Estrin. The Impact of Routing Policy on Internet Paths. To appear: *Proceedings of IEEE INFOCOM'2001*, April 2001.

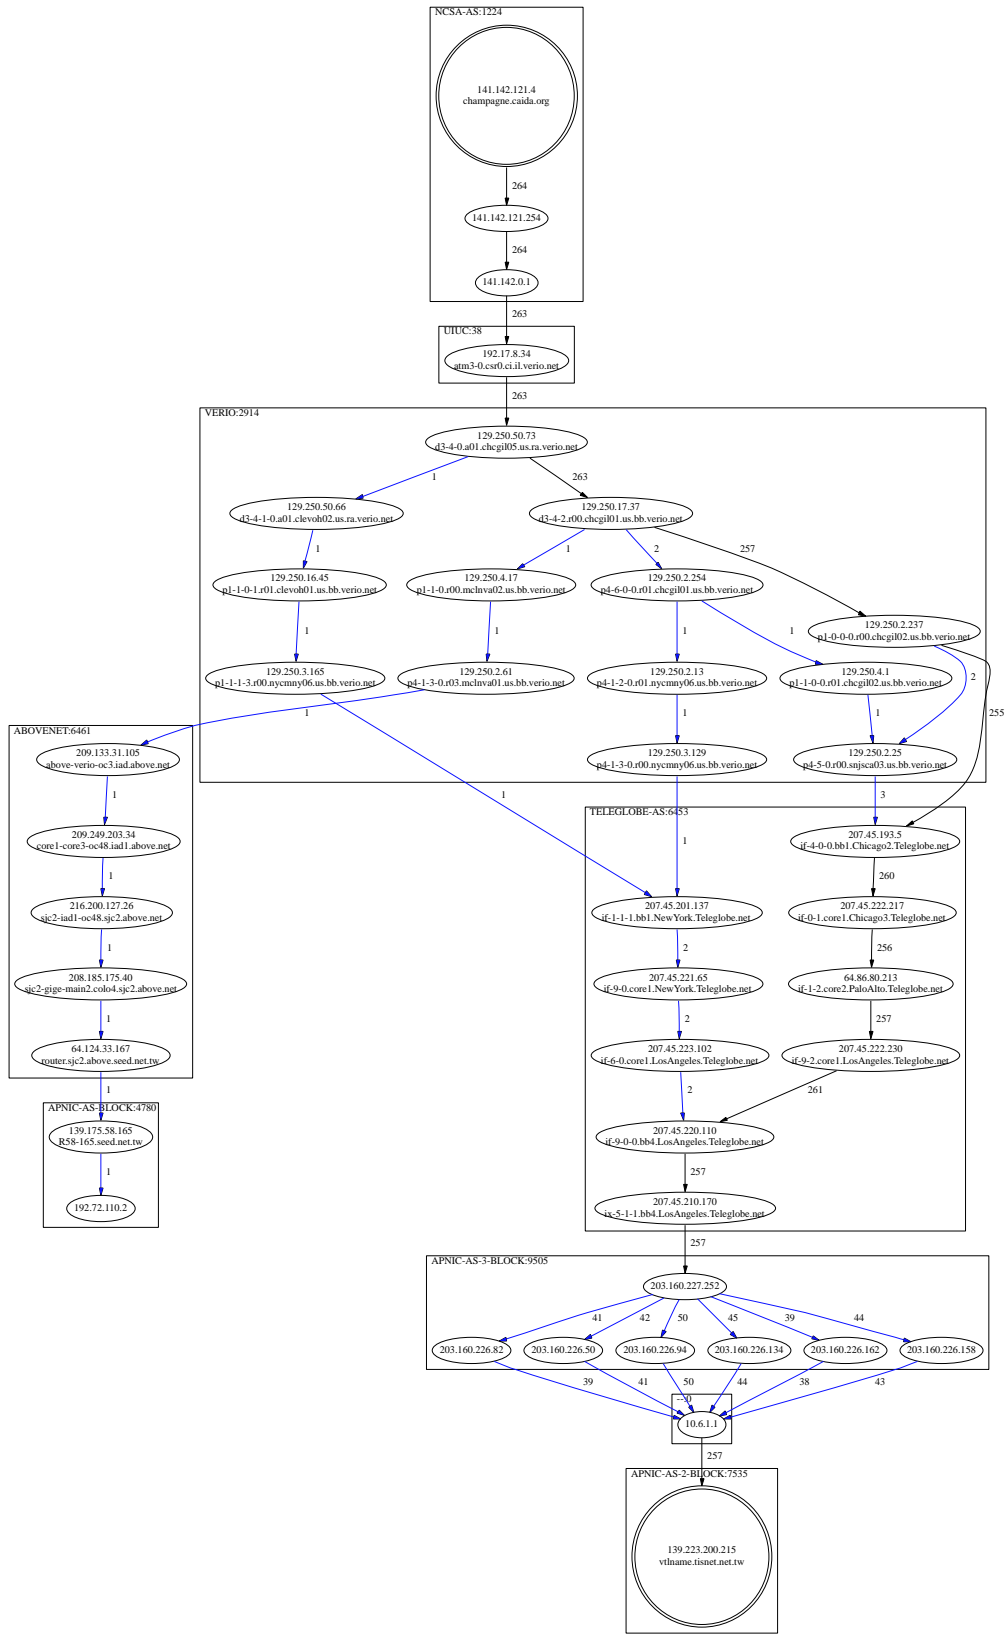


Figure 9: An example of Skitter data referenced against AS domain data provided by CAIDA. Notice the well-balanced six-edged diamond at bottom.