

Analysis of Overlay-Underlay Topology Correlation using Visualization *

Vinay Anja Marco Robert Dorothea
Aggarwal Feldmann Gaertler Görke Wagner
TU München Universität Karlsruhe

Abstract

Taking the physical Internet at the Autonomous System (AS) level as an instance of a complex network, and Gnutella as a popular peer-to-peer application running on top of it, we investigated the correlation of overlay networks with their underlying topology using visualization. We find that while overlay networks create arbitrary topologies, they differ from randomly generated networks, and there is a correlation with the underlying network. In addition, we successfully validated the applicability of our visualization technique for AS topologies by comparing Routeviews [10] data sets with DIMES [5] data sets, and by analyzing the temporal evolution in the Routeviews data sets.

1 Introduction

The contemporary Internet can be defined as a collection of segregated routing domains called Autonomous Systems (AS), each having their own administration and independent routing policies. Routing information between ASes is exchanged via an exterior gateway protocol such as BGP. The AS network possesses an implicit hierarchical structure where the ASes can be categorized into backbones, national, regional or local providers, and customers. The graph of the ASes, where nodes represent different ASes, and edges correspond to traffic trade agreements between the ASes, is an instance of a complex system. The widely popular peer-to-peer (P2P) file sharing systems create overlay topologies on top of the AS graph. Recent investigations have shown that overlay topology does not

*This work is partially supported by the EC FET Open Project DELIS, and by the DFG grant WA 654/12-1,2.

appear to be correlated with the underlying network topology [2]. This has potential performance drawbacks, as communication paths between P2P neighbors are often not optimal. We aim to enhance our understanding of the overlay-underlay interactions using visualization and the concept of cores [11, 3]. The k -core of an undirected graph is defined as the unique subgraph obtained by recursively removing all nodes of degree less than k . A node has *coreness* ℓ , if it belongs to the ℓ -core but not to the $(\ell + 1)$ -core. The ℓ -core layer is the collection of all nodes having coreness ℓ . The *core* of a graph is the k -core such that the $(k + 1)$ -core is empty. In general, the core decomposition of a graph results in disconnected subgraphs but in the case of the AS graph, we observe that all k -cores stay connected, a nice feature. Cores have been frequently used for network analysis [7, 6].

The paper is organized as follows. In Section 2, we present an analysis of overlay-underlay topology correlation using visualization. To reduce the bias in our visualization technique and to lessen the presence of artifacts in the AS graph analysis, we compare the Routeviews and DIMES data sets in Section 3. This is followed by an analysis of the temporal evolution of the AS graph in Section 4, and a summarization of our results in Section 5.

2 Evaluating Overlay Networks

In recent times, the growing share of P2P traffic in the Internet has generated a lot of interest in evaluating overlay networks. Overlays are formed at the application layer, but the actual data flow takes place at the network layer. Since the neighborhood selection process of overlay networks is largely arbitrary, it becomes interesting to analyze how much the neighborhood selection process of P2P protocols respects the underlying Internet topology.

Due to its high popularity [12] and open-source design, the Gnutella [8] network was chosen for our investigation. Using a combination of active and passive measurement techniques, we collected a large number of Gnutella node pairs (IP addresses) which were direct neighbors at some point of time. We call such a node pair denoting Gnutella neighbors as a Gnutella edge. The data was gathered by instrumenting a popular Gnutella vendor software GTK-Gnutella [9], and allowing it to crawl the Gnutella network non-intrusively. In order to reduce the bias in the measurement data, we concentrated on gathering remote Gnutella edges, i. e., Gnutella node pairs where none of the nodes represents our crawler. The investigation methodology and the analysis results are explained in detail in [2].

Among other things, we find that nodes in overlay networks do not seem to significantly bias their neighborhood choices towards network proximity. This means that P2P nodes form neighborhoods irrespective of the network layer connectivity. Since the amount of traffic involved in P2P file-sharing systems is huge, this has

potential performance drawbacks. Namely, data often flows through longer, sub-optimal paths for no apparent necessity. As an example, a node in Hamburg may access files from New York, which may also be available at a node in Frankfurt. This implies that biasing the neighborhood selection in P2P systems at the application layer to respect network connectivity can lead to performance improvements. This is the eventual goal of our investigation. To gain a better understanding of the overlay-underlay interactions, we investigate this using visualization.

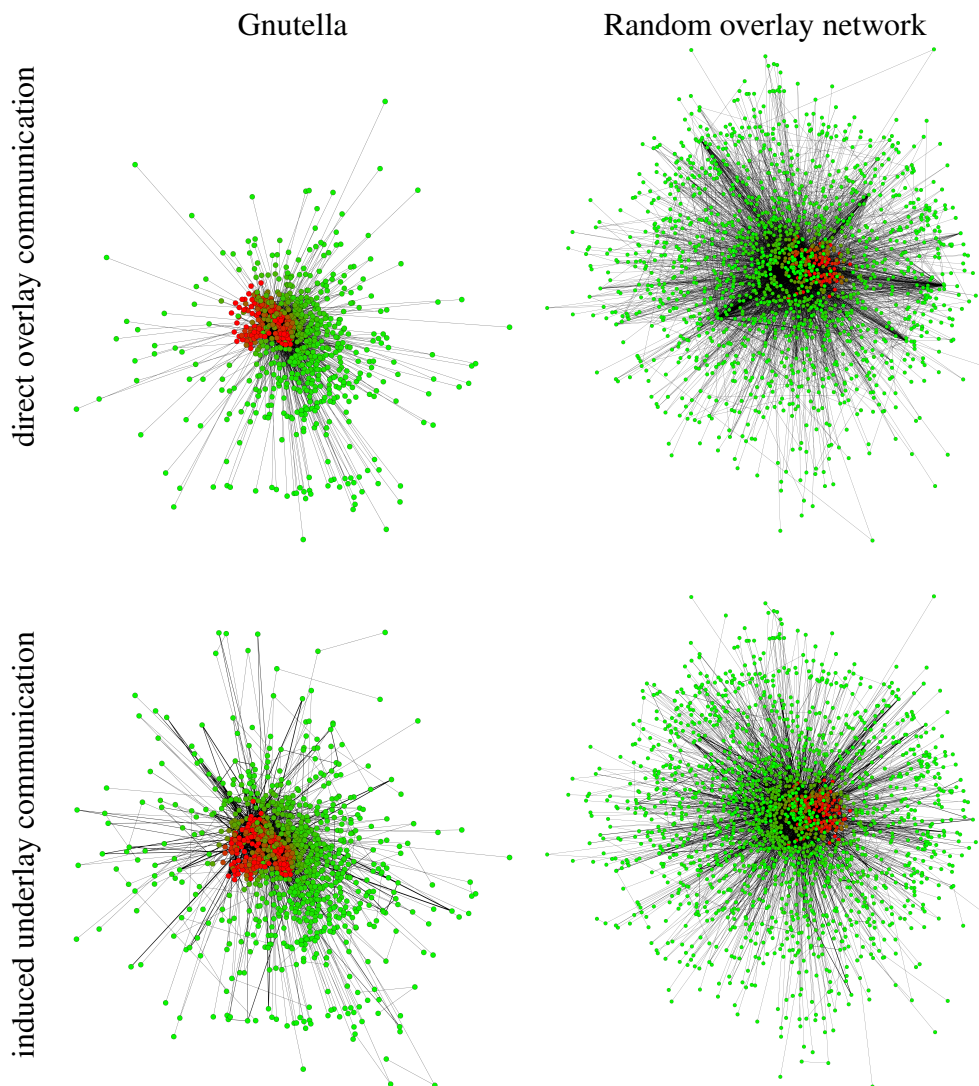


Figure 1: Comparison of overlay communication in a real (Gnutella) and a generated (random) network, using visualization [4]

Employing the visualization technique presented in [4], we first established

a layout for the underlying AS graph. Using the previously mentioned Gnutella measurement setup [2], we collected Gnutella logs for one week in April 2004, and mapped the IP addresses of Gnutella nodes to ASes using the BGP table dumps offered by Routeviews [10] during the week of April 14, 2004. We then compare the Gnutella peering connections with random peerings, where the communication end-points are uniformly randomly drawn from the whole IP space and then mapped to ASes. For the Gnutella graph, we obtained 2964 unique edges (direct P2P connections) corresponding to 754 ASes, while for the random network, we generated 4975 unique edges which corresponded to 2095 ASes. We ignored P2P edges within the same AS. In order to highlight the corresponding communication paths, we consider two models: *direct overlay communication* and *induced underlay communication*.

In the direct overlay communication model, we consider an AS-abstracted view of the direct P2P communication graph, where nodes are ASes and an edge connects two nodes if there exists a direct P2P communication between the corresponding ASes. The *appearance weight* of an edge is the number of such communications between the corresponding ASes. Note that the edges in this model disregard the underlying topology.

In the case of the induced underlay communication model, we associate each overlay edge with the corresponding underlay path using AS path information from Routeviews data sets. This is one possible path the P2P communication may traverse. We remove all edges of the underlay graph that are not used in any overlay communication. The *appearance weight* of an edge denotes the number of paths it appears in. This refines the meaning of the edges in the original underlay network, i. e., an edge is present between two ASes if and only if they have a traffic exchange agreement and a P2P communication is routed through it.

Figure 1 shows the corresponding visualizations. The random overlay network has several nodes with large degree which are located in the periphery, while Gnutella has almost none.

The different characteristics of Gnutella and random networks in the two models motivate us to investigate structural dependencies between the induced underlay communication model and the actual underlay network itself. Edges in the underlay network are not equally loaded as some edges appear in more communication paths than others. As it is not possible to measure the actual traffic on the individual edges, we consider a simplified model where a single communication causes one unit of traffic to be routed. The appearance weight of an edge in the underlay communication model thus corresponds to its load. The real load of an edge in the underlay network (including all the traffic caused by other applications) is naturally larger. Comparing these two loads reveals whether the P2P communication has characteristics similar to the accumulated load. This helps in understanding and enhancing the underlay network topology and application level

routing techniques. However, measuring the traffic load in the underlay network is not trivial. Even in a simplified model where we consider the load to be equal to the number of appearances in router-path announcements, the measurement heavily depends on the collection of routers and the collection process and is thus biased. Hence, we compare the appearance weight with node-structural properties of the corresponding end-nodes in the original underlying network. We focus on the properties degree and coreness, as both have been successfully applied for the extraction of customer-provider relationship as well as visualization [1, 6], as these properties reflect the importance of ASes. We establish the relation between the edge weight and the node structural measurements by systematically comparing the weight with both the smaller and the larger end-node value. Figure 2 shows the plots of the weights versus the degree.

From the plots it is apparent that the appearance of an edge and its end-nodes' degrees are not correlated in both the Gnutella and the random network, as no pattern is observable. Also, the distributions are similar as the majority of edges are located in the periphery of the network where the maximum degree of the end-nodes is small. We thus hypothesize that the relation of load in the P2P network and node degree in the underlying network is the same in both the Gnutella and the random network.

However, the situation changes when considering the coreness instead of the degree. From Figure 3, we can observe that although there is no correlation in any of the two networks, the distributions are different. In the random network the distributions are very uniform, which is a reflection of its random nature. In the case of Gnutella almost no heavy edge is incident to a node with small coreness, as can be seen in the minimum-coreness diagram. Positively speaking, most edges with large weights are incident to nodes with large minimum coreness. Interpreting coreness as importance of an AS, these edges are located in the backbone of the Internet. The same diagram for the random network does not yield a similar significant distribution, thus denying a comparable interpretation. For instance, in the random network, there exist edges located in the periphery that are heavily loaded. As an aside, backbone edges need not necessarily be heavily loaded in either network.

These observations lead us to conclude that the Gnutella network differs from random networks and appears to be correlated to the underlay network. Finding a sound reasoning for this correlation is part of our ongoing research.

3 Comparing Different Data Sources

The Oregon Routeviews Project is one of the major repositories for snapshots of the AS network using looking glasses. In contrast, the DIMES [5] project

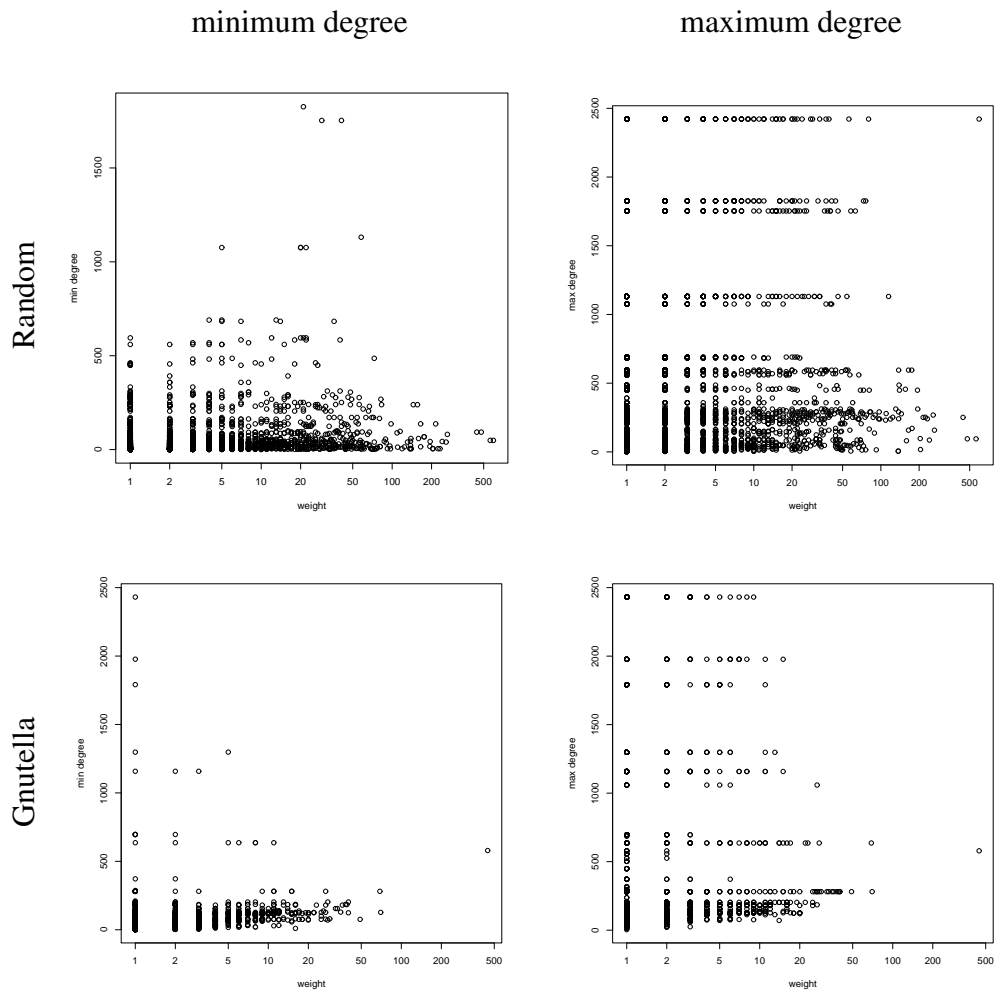


Figure 2: Comparing appearance weight with minimum and maximum degree of the corresponding end-nodes in Gnutella and the random network. Each data point represents an edge, the x -axis denotes the appearance weight and the y -axis reflects the degrees of the end-nodes. All axes use logarithmic scale.

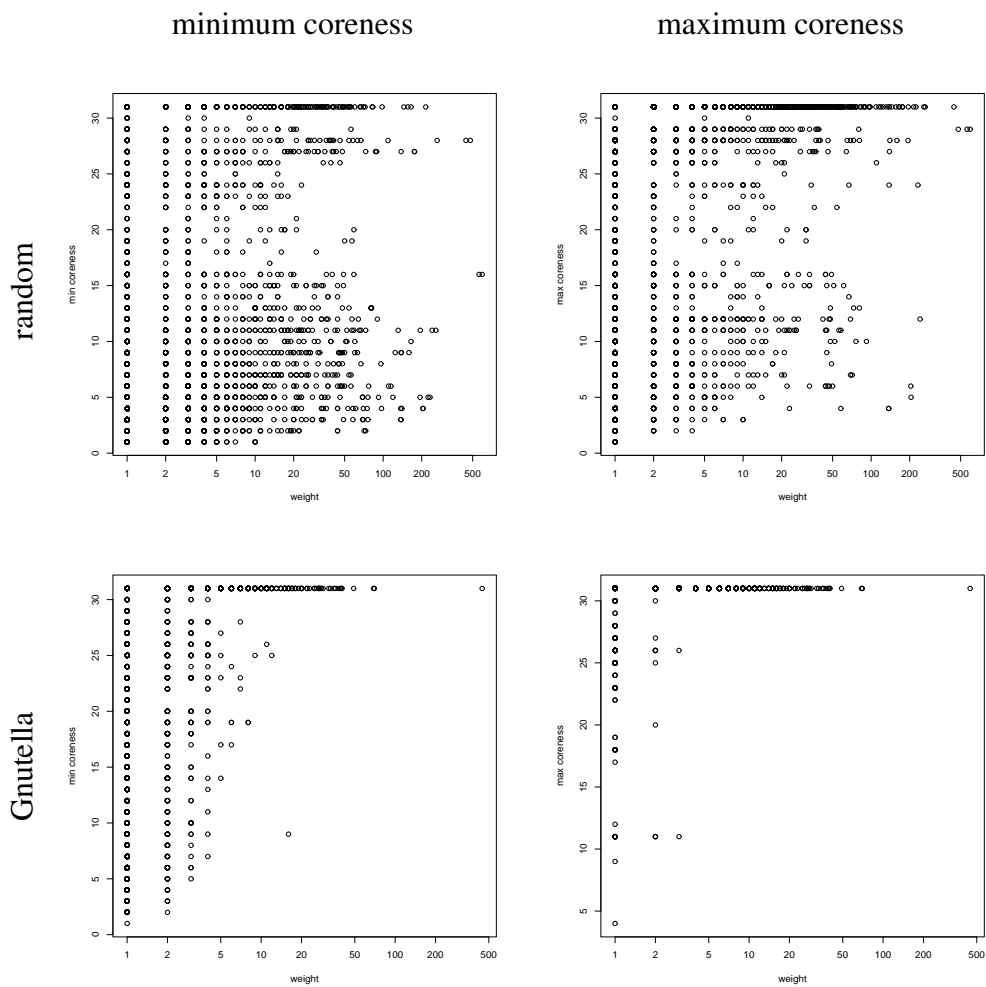



Figure 3: Comparing appearance weight with minimum and maximum coreness of the corresponding end-nodes in Gnutella and the random network. Each data point represents an edge, the x -axis denotes the appearance weight and the y -axis reflects the coreness of the end-nodes. All axes use logarithmic scale.

extracts AS relations using traceroute experiments. Since the structure of the AS network used in Section 2 was obtained using Routeviews data, we compare here the Routeviews AS topological map with that of DIMES. A similar structure of topological maps obtained from these two different techniques implies that our analysis is not biased by the source of AS data.

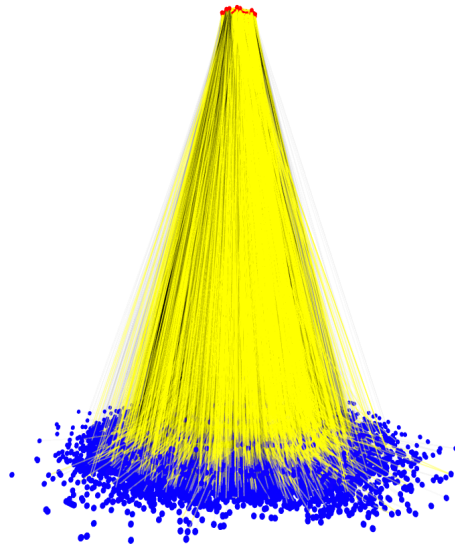
Figure 4 visualizes the edges obtained from Routeviews and DIMES data sets, where an edge denotes a traffic exchange between two ASes, represented by the end-nodes of the edge. In order to highlight structural properties, we colored edges that were only obtained from Routeviews as cyan, edges that were only obtained from DIMES as yellow, and edges that were present in both data sets as black. We used the visualization technique [4] on the union of the data sets in order to obtain the global layout. In Figure 4(a) and 4(b) only the parts obtained from the individual data sources are displayed. Figure 4(c) and 4(d) show union and intersection of the two data sets respectively. The data sets correspond to the period of March to June 2005. We obtained 48,073 edges (corresponding to 20,406 ASes) from Routeviews and 38,928 edges (corresponding to 14,154 ASes) from DIMES. Of these, 21,725 edges belong exclusively to Routeviews, and 12,580 edges exclusively to DIMES. The rest of the edges are common to both data sets. The union of the two data sets thus results in 60,653 unique edges (corresponding to 20,612 ASes).

A first glance at Figure 4 shows that while the visualizations are very similar, the DIMES data set is slightly smaller and the geometric difference of the set of collected edges of Routeviews and DIMES is surprisingly large (about 42% overlap). In other words, 58% of the edges appear in only one data set. An interesting observation is that many edges that are only discovered by DIMES are incident to the core.

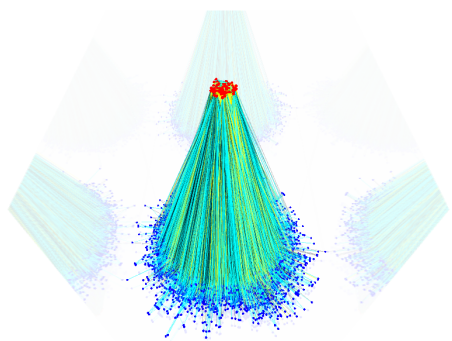
Figure 5 shows the plots of the coreness of the end-nodes of the edges (which represent ASes) versus their rank, positioned in the non-decreasing sorted sequence. The coreness is calculated in the graph that consists of the union of the two data samples. This enables us to set up a less biased comparison. The Routeviews data sample is plotted as a solid line, while the DIMES sample is dotted. Figure 5(a) plots a data point for each edge belonging to Routeviews or DIMES using the maximum coreness of the end-nodes (as y -axis), while Figure 5(b) shows the same scenario using the minimum coreness. A similar comparison is made in Figures 5(c) and 5(d) where the common edges are omitted. Thus the solid lines represent the distribution of edges that are exclusively observed by Routeviews, and the dotted lines correspond to the exclusive part of the DIMES sample. In principle, the distributions of Routeviews and DIMES are very similar, except for the broad tail of the Routeviews distribution observed in Figure 5(c). However the overall similarity of the plots together with the resembling visualizations in Figure 4 reveal that the Routeviews and DIMES data sets are indeed similar.




(a) Routeviews



(b) DIMES




(c) Union of Routeviews and DIMES




(d) Intersection of Routeviews and DIMES


Figure 4: Visualizations of AS networks obtained from Routeviews and DIMES using k-core technique [4]. Colour code: cyan-Routeviews, yellow-DIMES, black-both Routeviews and DIMES




(a) maximum coreness



(b) minimum coreness



(c) maximum coreness



(d) minimum coreness

Figure 5: Comparison of coreness distributions of the edges. Figure (a) and (b) compare Routeviews (solid) with DIMES (dotted), while Figure (c) and (d) compare the exclusive sets. The x-axis denotes the number of edges, and y-axis the minimum or maximum coreness.

4 Macroscopic Evolution

To ensure that our analysis of the AS graph structure is not biased by the time of measurement, we analyze the temporal evolution of the AS graph over a longer period of time. We use the graph-theoretical concept of k -cores to track the general shape of the AS network over time. Again, the visualization technique [4] relates the coreness of an AS to its position in the layout very well (recall that nodes with large coreness are placed in the center while nodes with small coreness are placed in the periphery) as is illustrated in Figure 6.




Figure 6: Visualization of the AS network (Jan 1, 2005) using [4]. Small blue nodes have small coreness while big red nodes have large coreness.

We observe that during the period of April 2001 to April 2005, the number of nodes in the AS graph increases by about 2000 nodes per year, the number of edges increases by 4800 edges per year and the maximum core number has increased from 18 to 26. Although the network increases in absolute terms and especially, the individual core levels grow, their relative sizes remain stable. Similar to the rings of a tree trunk, Figure 7 illustrates the temporal evolution of the relative proportions of the k -shells, i. e., collection of nodes with coreness k . In this figure, the thickness of one strip corresponds to the fraction of nodes that have a

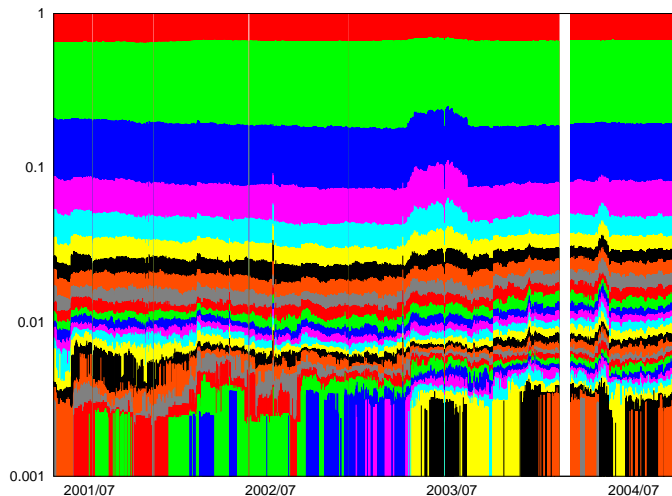


Figure 7: Relative size of cores.

given coreness. The lowest strip represents the maximum core while the highest strip reflects the 1-shell. Besides the stability of k -shells with $k \leq 15$, it is also observable that the size and coreness of the maximum core increases over time. The growth in the coreness is not monotonic and has big fluctuations. White vertical strips indicate absence of data due to technical problems in the external data collection process.

Furthermore, the relative distances of the ASes to the ‘center’ in the visualizations remains roughly the same. Figure 8 represents the two-dimensional distributions of nodes in the layout which is given in Figure 6. This distribution is obtained by superimposing a grid onto the layout and counting the number of nodes in the individual grid cells. It clearly illustrates the general volcano-like shape. In contrast to previous attempts employing techniques like force-directed or spectral methods where the majority of nodes were placed at a central position, here only 6-10% of the nodes are placed close to the center, constituting a peak, and the majority of the nodes, i. e., those with coreness two or three, are placed in a concentric annulus around it. Using several snapshots over time, we found a positive correlation of $0.67 - 0.78$ between the distance from the center and the coreness. The significance of this correlation is increased by the following facts: first, nodes in the 1-core are placed very close to their anchor nodes in higher core-levels which can be quite scattered or close to the center, and second, nodes with coreness two or three constitute the majority of nodes and occupy a broad annulus rather than a ring. Figure 9 shows several snapshots of the AS networks where only nodes with coreness two (green) and three (red) are drawn. The whole layout was again computed using the technique of [4]. It is apparent that the an-

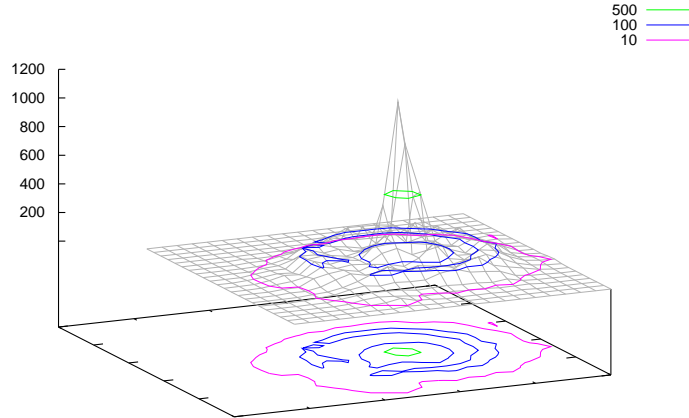


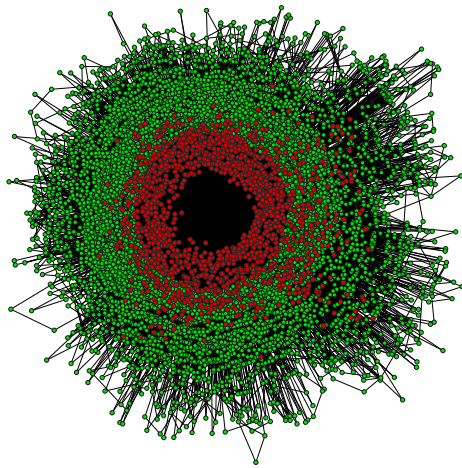
Figure 8: Histogram of the data points of Figure 6 showing the 2-dimensional distribution of nodes in the layout, showing only 6-10% of nodes close to center.

nuli remain constant over time. Except for a few nodes, most nodes of the two- and three-shell are placed in the periphery of the layout which reflects their role in the network.


On the whole the layouts for the different snapshots of the AS graph over four years reveal the same properties, thus implying that our overlay analysis is not biased by the time of measurement. It also verifies the usability of the visualization technique [4] to visualize and analyze the AS network.

5 Conclusion


Using visualization and concept of cores, we established that while overlay networks use an arbitrary neighborhood selection process, their topology differs from randomly generated networks. There exists some correlation between the overlay and the underlay network topologies. By comparing the Routeviews and DIMES data sets, we confirm that different data sources or collection processes do not affect our view of the AS graph. The analysis of the temporal evolution of the Routeviews data sets helps us to verify the global stability of the AS graph over an extended period of time. Apart from showing that our overlay-underlay analysis is not biased by the source of AS data or the time of measurement, this also demonstrates the applicability of the visualization technique [4] to analyze multiple sources of AS data over long time periods.




(a) 2002



(b) 2003



(c) 2004



(d) 2005

Figure 9: Snapshots of the AS network (2002-2005) where only the second (green) and the third (red) shell is drawn, using k-core technique [4].

References

- [1] S. Agarwal, L. Subramanian, J. Rexford, and R. Katz. Characterizing the Internet Hierarchy from Multiple Vantage Points. In *Proceedings of IEEE Infocom 2002*, 2002.
- [2] Vinay Aggarwal, Stefan Bender, Anja Feldmann, and Arne Wichmann. Methodology for Estimating Network Distances of Gnutella Neighbors. In *GI Jahrestagung*, volume 2, 2004.
- [3] V. Batagelj and M. Zaveršnik. Generalized cores. Preprint 799, Universtiy of Ljibljana, 2002.
- [4] Michael Baur, Ulrik Brandes, Marco Gaertler, and Dorothea Wagner. Drawing the AS Graph in 2.5 Dimensions. In *Proceedings of the 12th International Symposium on Graph Drawing (GD'04)*, volume 3383 of *Lecture Notes in Computer Science*, pages 43–48, 2005.
- [5] DIMES – Distributed Internet MEasurements & Simulations. <http://www.netdimes.org/>.
- [6] Marco Gaertler and Maurizio Patrignani. Dynamic Analysis of the Autonomous System Graph. In *IPS 2004 – Inter-Domain Performance and Simulation*, pages 13–24, March 2004.
- [7] Christos Gkantsidi, Milena Mihail, and Ellen Zegura. Spectral Analysis of Internet Topologies. In *IEEE Infocom 2003*, 2003.
- [8] Gnutella v0.6 rfc. <http://www.the-gdf.org/>.
- [9] GTK-Gnutella. <http://gtk-gnutella.sourceforge.net/>.
- [10] University of Oregon Routeviews Project. <http://www.routeviews.org/>.
- [11] S. B. Seidman. Network structure and minimum degree. *Social Networks*, 5:269–287, 1983.
- [12] Slyck. <http://www.slyck.com/>.