A CLUSTERING APPROACH FOR EXPLORING THE INTERNET STRUCTURE

Avishai Wool, Gonen Sagie

School of Electrical Engineering,
Tel Aviv University,
Ramat Aviv 69978, Israel.
yash@acm.com, gonen@kagoor.com

ABSTRACT

This paper proposes several clustering algorithms to explore the topology of the AS-graph. Using these algorithms, we are able to view the Internet topology, as encapsulated by the AS-graph, at three levels of abstraction: (1) a low-resolution view, which shows a coarse cluster cover of the whole AS-graph, (2) a mid-resolution view, showing the relationships between dense cores inside the coarse clusters, and (3) a high-resolution view of individual high-density cores.

Using these tools we discovered that the Internet's central Dense Core is at least twice as large as what was previously reported. Attached to the central Dense Core, in what is known as the Transit Core, our clustering tools discovered several clusters, which we call the Regional Dense Cores: Interestingly, these clusters are as dense as the central core itself. Our tools highlight the fact that the Regional Dense Cores are mostly connected to the central Dense Core and less well connected among themselves.

1. INTRODUCTION

1.1. Background and Motivation

The connectivity of the Internet crucially depends on the relationships between thousands of Autonomous Systems (ASes) that exchange routing information using the Border Gateway Protocol (BGP). These relationships can be modeled as a graph, called the AS-graph, in which the vertices model the ASes, and the edges model the peering arrangements between the ASes.

Significant progress has been made in the study of the AS-graph's topology over the last few years. Several characteristics of the AS-graph have been discovered empirically (cf. [1]). Most of these properties deal with the distribution of vertex degrees (i.e., the number of peers that an AS has), and the discovery that this distribution is heavy-tailed and obeys so-called power-laws. Less progress has been made in the areas of obtaining better conceptual models of the AS-graph, exploring the graph at different levels of abstraction and detail, and visualizing the AS-graph.

Therefore, the goal of our work has been threefold:

1. To devise tools that allow us to explore the Internet structural topology, as encapsulated by the AS-graph.
2. To use these tools to gain a better understanding of the Internet connectivity.
3. To produce aesthetically-pleasing Internet drawings, which highlight interesting features of the AS-graph.

1.2. Related Work

1.2.1. Internet Power-Law Models

In a ground-breaking paper, Faloutsos et al. [1] discovered that several parameters of the Internet topology are governed by power-laws. The power-laws describe skewed distributions of the graph properties, including the vertex degree. They showed that these power-laws hold for three Internet snap-shots taken between 1997 and 1998.

Barabási and Albert [2] introduced a topology model (the BA model) for diverse generic networks. The model is based on two mechanisms: (i) networks expand continuously by the addition of new vertices, and (ii) new vertices attach preferentially to sites that are already well connected. They showed that these two mechanisms suffice to produce networks that are governed by the power-laws, similar to the laws discovered by [1] in the AS connectivity graph. This model was later refined in [3] to allow adding links to existing edges, and to rewire existing links.

1.2.2. Conceptual Models

Tauro et al. [4] suggested the "jellyfish" model, a conceptual model for the Internet topology. They introduced metrics to qualify the topological importance and significance of the nodes, and used these metrics to a define topological model. They argued that the topology resembles a jellyfish where the Internet core corresponds to the middle of the cap, which is surrounded by many "tentacles".

Subramanian et al. [5] suggested a 5-tier hierarchical layering of the AS-graph. Their top tier, called the Dense Core, is defined as a subset of ASes whose edge density[1] is $> 50\%$. Their tiering agrees with the jellyfish model of [4] in that, implicitly, they assume a single dense core, which is surrounded by the Transit Core, and then by the Outer Core etc. Using a simple greedy algorithm they identified a Dense Core of 20 ASes.

There are several synthetic Internet topology generators available for use, such as BRITE [6] and Inet [7, 8]. Both mostly use degree-based methods that attempt to match the empirical power-law degree distribution. A critique of pure degree-based network generators appears in [9], which claims that such synthetic networks mis-represent hierarchical features of the Internet structure. Bu and Towsley [10] find that degree-based generators differ significantly in their clustering coefficients. Their work proposes an alternative degree-based generator that more closely matches the clustering behavior of the measured AS-graph.

[1]The density $\varrho(k)$ of a subgraph with k vertices is the fraction of the $k(k-1)/2$ possible edges that exist in the subgraph.

Based on the work reported here, we have recently shown [11] that the BA model, and its implementations in BRITE and Inet, are are fundamentally *unable* to produce synthetic topologies with a substantial dense core. Bar, Gonen, and Wool [12] recently proposed an improved BA-type model, which addresses the discrepancies of the BA model with respect to the lack of leaves, the value of the power law parameter γ, and the lack of a dense core.

1.2.3. *Drawing the Internet*

Several efforts have been made to visualize the Internet topology. Cheswick, Burch and Branigan [13, 14] studied the Internet topology and drew Internet maps at the router level. They produced drawings of the whole Internet which illustrated what the Internet looks like from a single host's point of view. Their drawings only show reachability tree graphs and not a complete map.

CAIDA's Skitter project [15, 16] describes a visualization that shows a macroscopic snapshot of the Internet core. Skitter visualizes interconnection relations between ASes correlated to their geographical location and connectivity level. Skitter emphasizes global geography (longitude) but does not provide the ability to explore different levels of granularity.

1.3. Contributions

The main theme in our work is *clustering*: the ability to aggregate ASes into meaningful groups. This paper proposes several clustering algorithms to explore the topology of the AS-graph. Using these algorithms, we are able to view the Internet topology, as encapsulated by the AS-graph, at three levels of abstraction: (1) a low-resolution view, which shows a coarse cluster cover of the whole AS-graph, (2) a mid-resolution view, showing the relationships between dense cores inside the coarse clusters, and (3) a high-resolution view of individual high-density cores.

Using these tools we discovered that the Internet's central Dense Core is at least twice as large as what was previously reported. Attached to the central Dense Core, in what is known as the Transit Core, our clustering tools discovered several clusters, which we call the Regional Dense Cores: Interestingly, these clusters are as dense as the central core itself. Our tools highlight the fact that the Regional Dense Cores are mostly connected to the central Dense Core and less well connected among themselves.

An important feature of our clustering algorithms is that they allow us to produce drawings of the AS-graph topology at different levels of abstraction. These drawings let us see all the above-mentioned features in a clear and explicit way.

2. TWO-HOP CLUSTERING

2.1. The Algorithm

To obtain a coarse clustering of the whole AS-graph, we used a variant of sparse-partition clustering (cf. [17, 18]), which we call Two-Hop Clustering (*2HC*). Using this algorithm we identify a giant cluster, all of whose members are at most two hops away from the cluster's center (which is UUnet, AS 701, the highest-degree AS). However, the clustering algorithm also shows the existence of fairly high-degree nodes with large clusters around them outside this giant cluster. We visualized this clustering to obtain a low-resolution, birds-eye, view of the AS-graph structure.

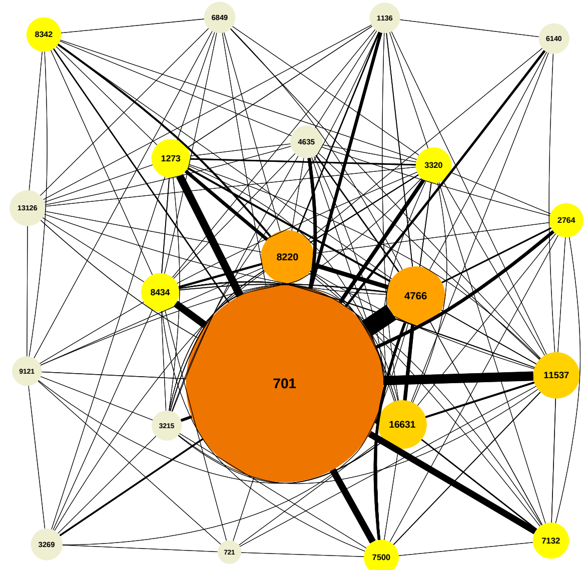

Fig. 1. A visualization of the *2HC* super-graph we found in the AS-graph (graph layout by `neato` [19]). The graph contains 21 clusters connected by 114 edges, covering 12,553 ASes. AS 701 is the main AS for UUnet—the highest degree AS. The radius of the vertex $S_i$is proportional to $\sqrt{|S_i|}$, so the circle area is proportional to the number of ASes in the cluster (for small clusters the label size may make the circle larger than it should be). The edge width is proportional to the number of the joint vertices between the clusters.

2.2. Applying *2HC* to the AS-graph

We implemented the *2HC* algorithm and applied it to the AS-graph. We set $Min\_Cluster\_Size = 100$ to discard small clusters. A visualization of the super-graph $\widetilde{G}$ is shown in Fig. 1. The super-graph $\widetilde{G}$ contains 21 clusters connected by 114 edges, covering 12,553 ASes. The remaining 2,429 ASes are not organized into significant clusters and are ignored.

Based on the "Jellyfish" analogy of [4], we expected the algorithm to produce a giant cluster including the Internet core, followed by very small clusters (if any) representing the fragments of the "tentacles" disconnected from the core.

Interestingly, the *2HC* algorithm results were not exactly what we expected. The first cluster is indeed a giant cluster, with 9,924 nodes and 24,241 internal edges. However, after removing the interior of the giant cluster (comprised of 2,648 nodes and 6,366 edges) from the AS-graph, *2HC* still found three more clusters with more than 500 nodes each, and a total number of 21 clusters with over 100 nodes, see Fig. 1.

3. LOOKING INTO THE GIANT CLUSTER

To explore the structure inside the giant cluster, we used an approach that is based on dense k-subgraphs (DkS). We adapt parts of a theoretically-interesting DkS approximation algorithm [20], and show a practical clustering procedure using the DkS algorithm as a building block. This DkS-based clustering allows us to explore the structure of the Internet's core. We visualized the relationships among the DkS clusters as our mid-resolution view, and visualize individual dense clusters as the high-resolution view.

3.1. The Dense k-Subgraph (DkS) Problem

The clustering approach we use to explore inside the *2HC* clusters is based on the notion of *dense subgraphs*. In this setting, one wishes to maximize two parameters: the number of nodes in the subgraph (quantified by the parameter k), and the edge density of the subgraph (what fraction of the $k(k-1)/2$ possible edges exist in the subgraph).

The Dense k-subgraph (DkS) maximization problem was studied by Feige, Kortsarz and Peleg [20]. In the variant they studied, the parameter k is fixed, and the goal is to find the densest subgraph over k vertices. Their algorithm is actually comprised of three separate algorithms, named $A1$, $A2$, and $A3$, each of which finds a candidate dense subgraph. We found that their methods, and in particular the A3 algorithm, worked quite well on the AS-graph data.

The A3 algorithm accepts the target subgraph size, k, as a parameter, and outputs the densest subgraph it can of the prescribed size. However, what we are really interested in are large, non-trivial, subgraphs, with a density higher than some threshold *Min_Density*. Thus, we designed a binary search algorithm called *A3Bin*, which uses $A3$ as a subroutine. *A3Bin* returns the single largest DkS cluster it is able to find, which has a density of at least *Min_Density*.

3.2. Applying *A3Bin* to the AS-graph

We implemented the *A3Bin* algorithm and ran it on the giant cluster that was found by the *2HC* algorithm. We used a minimal density of $Min\_Density = 70\%$, and a minimal cluster size of $Min\_k = 7$.

The first and largest DkS cluster we obtained consists of 43 ASes, with 637 of the possible 903 edges (density $= 70.5\%$). See Fig. 3 for a visualization of this cluster. This is a significant discovery in and of itself: Previously, Subramanian et al. [5] claimed that the Dense Core consists of only 20 ASes. Our algorithms discovered that the Dense Core is at least twice as large—and much denser (over 70% density). This indicates that the *A3Bin* algorithm, with its solid theoretical foundation, is clearly superior to the simple greedy heuristic used by [5].

We repeatedly applied *A3Bin* and built the super-graph $\widehat{G}$. The resulting super-graph has 5 clusters, of sizes 43, 14, 8, 8 and 7. See Fig. 2 for a visualization of $\widehat{G}$. In other words, there are multiple Dense Cores—an observation which does not fit a "jellyfish" conceptual model very well.

We are able to observe very clear geographic locality in the clusters. The largest DkS cluster consists mostly of US-based ASes, mixed with ASes from other regions. However, each of the four other clusters has a distinct regional flavor: *all* the ASes inside these clusters are headquartered in a single region. We found two European clusters, one in Canada and one in Australia/Japan. Therefore, we call these DkS clusters *Regional Dense Cores*.

Furthermore, the 5 DkS clusters are not all connected to each other. The Regional Cores are all connected to the worldwide 43-AS DkS cluster, which is reasonable since the giant *2HC* cluster is centered around AS 701—and AS 701 is a member of the 43-AS DkS cluster. However, only 2 of the 6 possible lateral Regional-to-Regional connections exist: the two European cores are directly

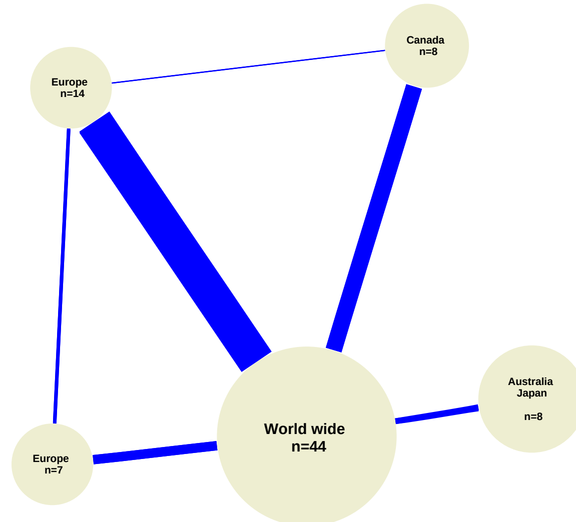

Fig. 2. A visualization of the DkS super-graph $\widehat{G}$ for the subgraphs we found inside the giant cluster. The graph consists of 5 clusters connected by 6 edges, covering 81 ASes. Each subgraph density is $\approx 70\%$. The vertex areas are proportional to the cluster size. The edge widths are proportional to the number of edges connecting the subgraphs.

connected to each other, and one European core is connected to the Canadian core. The Australian/Japanese core is connected only to the world wide cluster. This connectivity is typical of what [5] calls the Transit Core. Also, the structure of $\widehat{G}$ fits the pattern seen in the Skitter drawing [16], in which most connections appear between the (longitudes of) the US and Europe, or between the US and the Pacific rim. Note that in the Skitter drawing, one cannot distinguish between Canada and the US because they cover the same range of longitudes.

Full details of this paper can be found in [21].

4. REFERENCES

[1] G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos, "Power-laws and the AS-level Internet topology," *IEEE/ACM Trans. on Networking*, vol. 11, pp. 514–524, 2003.

[2] Albert-László Barabási and Réka Albert, "Emergence of scaling in random networks," *Science*, vol. 286, pp. 509–512, 15 October 1999.

[3] Réka Albert and Albert-László Barabási, "Topology of evolving networks: Local events and universality," *Physical Review Letters.*, vol. 85, no. 24, pp. 5234–5237, Dec. 2000.

[4] L. Tauro, C. Palmer, G. Siganos, and M. Faloutsos, "A simple conceptual model for Internet topology," in *IEEE Global Internet*, San Antonio, TX, Nov. 2001.

[5] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *Proc. IEEE INFOCOM'02*, New-York, NY, USA, Apr. 2002.

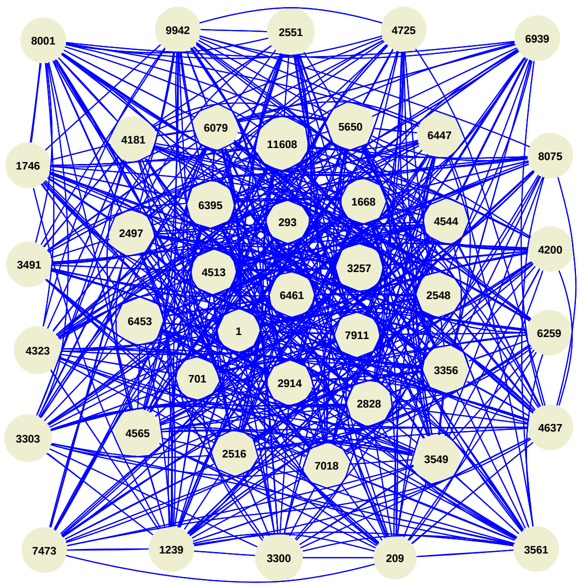

Fig. 3. 43-node subgraph of the Cluster 701 with 70.5 % density: 637 of the possible 903 edges are present.

[6] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRITE: An approach to universal topology generation," in *Proceedings of MASCOTS'01*, Aug. 2001.

[7] C. Jin, Q. Chen, and S. Jamin, "Inet: Internet topology generator," Tech. Rep. CSE-TR443-00, Department of EECS, University of Michigan, 2000.

[8] Jared Winick and Sugih Jamin, "Inet-3.0: Internet topology generator," Tech. Rep. UM-CSE-TR-456-02, Department of EECS, University of Michigan, 2002.

[9] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network topology generators: Degree based vs. structural," in *Proc. ACM SIGCOMM*, 2002.

[10] T. Bu and D. Towsley, "On distinguishing between Internet power-law generators," in *Proc. IEEE INFOCOM'02*, New-York, NY, USA, Apr. 2002.

[11] G. Sagie and A. Wool, "A clustering-based comparison of Internet topology models," 2003, Preprint.

[12] S. Bar, M. Gonen, and A. Wool, "An incremental super-linear preferential Internet topology model," in *Proc. 5th Annual Passive & Active Measurement Workshop (PAM), LNCS 3015*, Antibes Juan-les-Pins, France, Apr. 2004, pp. 53–62, Springer-Verlag.

[13] H. Burch and W. R. Cheswick, "Mapping the Internet," *IEEE Computer*, vol. 32, no. 4, 1999.

[14] W. R. Cheswick, H. Burch, and S. Branigan, "Mapping and visualizing the Internet," in *Proc. USENIX Annual Technical Conference*, 2000, pp. 1–12.

[15] "Visualizing Internet topology at a macroscopic scale," `http://www.caida.org/analysis/topology/as_core_network/`, 2002.

[16] B. Huffaker, A. Broido, K. Claffy, M. Fomenkov, S. McCreary, D. Moore, and O. Jakubiec, "Skitter: AS Internet graph," `http://www.caida.org/analysis/topology/as_core_network/AS_Network.xml/`.

[17] B. Awerbuch and D. Peleg, "Sparse partitions," in *Proc. 31st IEEE Symp. Foundations of Computer Science (FOCS)*, 1990, pp. 503–513.

[18] David Peleg, *Distributed Computing, A Locality–Sensitive Approach*, SIAM, 2000.

[19] "`neato`. graphviz - open source graph drawing software," version 1.8.9, 2002, `http://www.research.att.com/sw/tools/graphviz/`.

[20] U. Feige, G. Kortsarz, and D. Peleg, "The dense k-subgraph problem," *Algorithmica*, vol. 29, no. 3, pp. 410–421, 2001.

[21] G. Sagie and A. Wool, "A clustering approach for exploring the Internet structure," Tech. Rep. EES2003-7, Dept. Electrical Engineering Systems, Tel Aviv University, 2003, Available from `http://www.eng.tau.ac.il/~yash/ees2003-7.ps`.