

A model of Internet topology using k -shell decomposition

Shai Carmi*[†], Shlomo Havlin*, Scott Kirkpatrick[‡], Yuval Shavit[§], and Eran Shir[§]

*Minerva Center and Department of Physics, Bar-Ilan University, Ramat Gan 52900, Israel; [†]Benin School of Engineering and Computer Science, Hebrew University, Jerusalem 91904, Israel; and [§]School of Electrical Engineering, Tel-Aviv University, Tel Aviv 69978, Israel

Edited by H. Eugene Stanley, Boston University, Boston, MA, and approved May 7, 2007 (received for review February 8, 2007)

We study a map of the Internet (at the autonomous systems level), by introducing and using the method of k -shell decomposition and the methods of percolation theory and fractal geometry, to find a model for the structure of the Internet. In particular, our analysis uses information on the connectivity of the network shells to separate, in a unique (no parameters) way, the Internet into three subcomponents: (i) a nucleus that is a small (≈ 100 nodes), very well connected globally distributed subgraph; (ii) a fractal subcomponent that is able to connect the bulk of the Internet without congesting the nucleus, with self-similar properties and critical exponents predicted from percolation theory; and (iii) dendrite-like structures, usually isolated nodes that are connected to the rest of the network through the nucleus only. We show that our method of decomposition is robust and provides insight into the underlying structure of the Internet and its functional consequences. Our approach of decomposing the network is general and also useful when studying other complex networks.

fractals | networks | percolation

The Internet has become a critical resource in our daily life. It still suffers from many inefficiencies and as such has become a vibrant research subject. Identifying the Internet's topology and its properties is a prerequisite to understanding its distributed, collaborative nature and the potential for building new services. For example, a better understanding of the Internet's structure is vital for integration of voice, data and video streams, point-to-point and point-to-many distribution of information, and assembling and searching all of the world's information. Much activity is presently focused on "disruptive research," which goes beyond today's routing algorithms and protocols to exploit underused links and discover greater available capacity. The measurement work that forms the basis of our study in this article has discovered much "dark matter," previously undisclosed links in the Internet. Our article will address the question of where this unused capacity lies and how it can be exploited. The Internet structure is also relevant to the discussion of the many problems facing the Internet today, such as the security threat posed by viruses, worms, spyware, and spam (ref. 1 and the Global Environment for Network Innovations, www.geni.net).

Various tools from statistical physics, like scaling theory, percolation, and fractal analysis, have been applied to better understand the Internet and other complex networks (2–9). In particular, the surprising finding of the Internet's power-law degree distribution (10) has encouraged many scientists to use the degree (the number of immediate neighbors of a node) as an indicator of the importance and role of each node. However, using the degree as an indicator of function can be misleading both when looking at a single node and when looking at a distribution. For example, it has been shown (11) that topologies with a very different structure can have the same degree distribution.

Instead of node degree, we will use the " k -shell" decomposition to assign a shell index to each node in the Internet. Although node degrees can range from one or two up to several thousands, we find that this procedure splits the network into

40–50 shells only, the precise number depending on the measurement details. k -shell decomposition is an old technique in graph theory (12) and has been used as a visualization tool for studying networks such as the Internet (13) (see Fig. 2 for a visualization of our data). It involves pruning the network down to those nodes with more than k neighbors, just as has been studied in physics under the rubric of "bootstrap percolation," in the very different environment of regular lattices (14). Other studies (15–18),[¶] have developed the theory of some of the statistical properties of k -shells in random networks.

In this article, we apply the k -shell decomposition on a map of the Internet at the autonomous system (AS) level (of $\approx 20,000$ nodes, a result of the DIMES project (ref. 19 and www.netdimes.org; see *Methods: Distributed Mapping of the Internet*) to decompose the network into components with distinct functional roles. Surfacing the distinct role each component plays will demonstrate how this method helps us understand the large-scale function of a network as complex as the Internet. Also, it may reveal evolutionary processes that control the growth of the Internet.^{||} We believe this tool and the structure we uncover give important information about the function of each node in the graph.

Results: Decomposing the Internet into Three Components

First, we decompose the network into its k -shells. We start by removing all nodes with one connection only (with their links), until no more such nodes remain, and assign them to the 1-shell. In the same manner, we recursively remove all nodes with degree 2 (or less), creating the 2-shell. We continue, increasing k until all nodes in the graph have been assigned to one of the shells. We name the highest shell index k_{\max} . The k -core is defined as the union of all shells with indices larger or equal to k . The k -crust is defined as the union of all shells with indices smaller or equal to k .

We then divide the nodes of the Internet into three groups:

1. All nodes in the k_{\max} -shell form the nucleus.
2. The rest of the nodes belong to the $(k_{\max} - 1)$ -crust. The nodes that belong to the largest connected component of this crust form the peer-connected component.
3. The other nodes of this crust, which belong to smaller clusters, form the isolated component.

Author contributions: S.H., S.K., Y.S., and E.S. designed research; S.C., Y.S., and E.S. performed research; S.H. and S.K. contributed new reagents/analytic tools; S.C. and S.K. analyzed data; and S.C. and S.K. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

Abbreviation: AS, autonomous system.

[†]To whom correspondence should be addressed. E-mail: scarmi@shoshi.ph.biu.ac.il.

[¶]Shalit, A., Solomon, S., Kirkpatrick, S., International Workshop on Aspects of Complexity and Its Applications, Sept. 23–25, 2002, Rome, Italy.

^{||}Preliminary results were presented at the International Center for Theoretical Physics Workshop on the Structure and Function of Complex Networks, May 23–27, 2005, Trieste, Italy.

© 2007 by The National Academy of Sciences of the USA

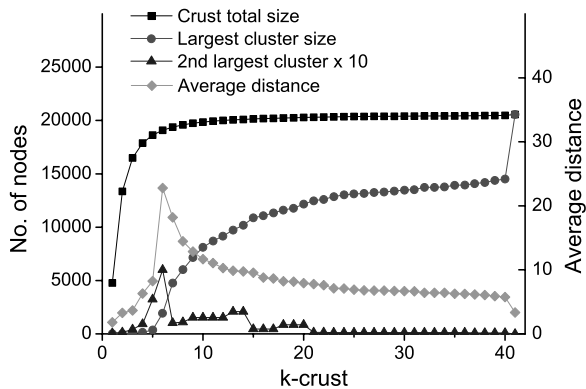


Fig. 1. For each k -crust, we plot the size of the crust (i.e., total number of nodes that belong to the crust), the size of the largest and second-largest connected components of the crusts (the second is magnified $\times 10$ to make it visible), and the average distance between nodes in the largest cluster of each crust.

We show in Fig. 1 how the sizes of the two largest components in each k -crust vary with the crust index. A percolation transition is apparent at $k = 6$. At this point, the size of the second largest cluster and the average distance between nodes in the largest cluster are sharply peaked (20, 21). This phase transition is similar to the transition found in ref. 22 when removing the high-degree nodes from a scale-free network. Above this point, the size of the largest cluster grows rapidly. At higher crusts it stabilizes, until it spans $\approx 70\%$ of the network at the $(k_{\max} - 1)$ -crust. When the nucleus is added, the network becomes completely connected. The jump in connectivity and dramatic decrease in the distances observed at the k_{\max} -shell (5.75 to 3.34) justifies our definition of it as the nucleus. However, even in the absence of the nucleus, 70% of the network remains connected (the peer-connected component). This connectivity offers important opportunities for transport control over the Internet. For example, to avoid congestion in the nucleus and increase total capacity, information can be sent by using only the more peripheral nodes of the peer-connected component. Nevertheless, a significant number ($\approx 30\%$ of the network) of other nodes are not connected in the $(k_{\max} - 1)$ -crust. Those nodes, which form the isolated component, are either leaves or form small clusters and can reach the rest of the network only through the nucleus. A schematic picture of the proposed decomposition (which we name a Medusa model for reasons explained below) is shown in Fig. 2.

Identifying the nodes that form the heart of the Internet, the nucleus, or “tier one,” is a problem that has been extensively investigated. For example, the nucleus might be defined as the set of all nodes with degree higher than some threshold. But this requires setting a free parameter, the degree threshold. Others (23) have defined it by using a growth process. Starting with an empty set, they add nodes to the nucleus in order of decreasing degree, retaining those for which the nucleus remains completely connected (a clique). We have found that heuristics to build up a maximal clique provide less accurate information about the network nucleus, for several reasons. First, they are not robust. Node degree is an ambiguous indicator of importance. If we consider other reasonable orderings of the nodes (for example, we ordered the nodes in descending order of their total number of links to nodes in the k_{\max} shell), the resulting clique differs in $>25\%$ of its constituent nodes. Moreover, it emphasizes American-based international carriers.

In contrast, our definition of the network’s nucleus is unique, parameter-free, robust, and easy to implement. Analyzing the ASes that are found in our proposed nucleus, we find that the set that participates is very stable over time. We repeated the construction by using data from 3-month intervals 3 and 6

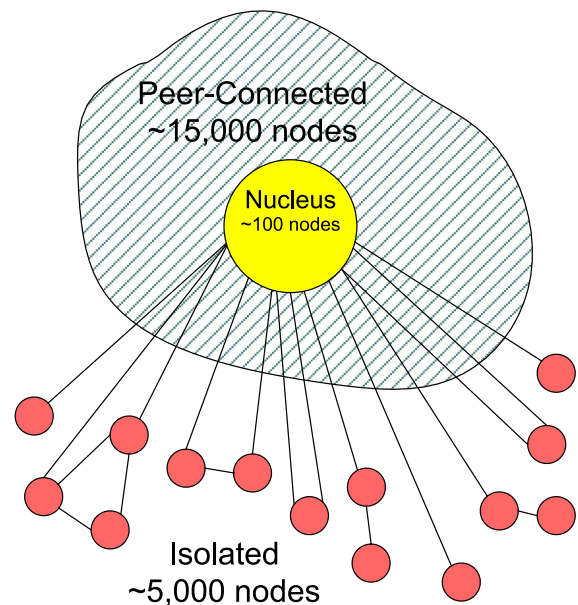



Fig. 2. Visualization of our data of the Internet at the AS level. (Upper) A plot of all nodes, ordered by their k -shell indices, using the program of ref. 13. The legend to the left denotes degree, and the legend to the right denotes k -shell index. (Lower) A schematic plot of the suggested Medusa model decomposition of the AS level Internet into three components.

months later than the data analyzed in this article and found changes consisting of a few percent of added sites and one or two sites that moved from the nucleus to a k -shell immediately before it. The actual ASes involved include all major intercontinental carriers (≈ 10 nodes), plus carriers and Internet exchange points equally distributed among countries in North America, Europe, and the Far East. The degree of nucleus sites ranged from $>2,500$ (ATT Worldnet) to as few as 50 carefully chosen neighbors, almost all within the nucleus (Google). The nucleus subgraph is redundantly connected, with diameter 2 and each node connected to $\approx 70\%$ of the other nucleus nodes, which provides k_{\max} -connectivity.

An interesting question arises: does the size of the nucleus increase with the Internet size and how? Although we have seen a steady increase in the size of the AS graph during the course of the DIMES project, we cannot yet separate the actual growth of the Internet from the increase in our measurement sensitivity. Thus, we are led to investigate random ensembles of scale-free networks, with parameters (such as degree distribution) similar to the real Internet (Fig. 3a). Note, that the random graph

k -crusts close to the percolation threshold show this behavior, up to some finite size effects (see Fig. 4*b*).

The isolated component includes nodes that are connected to the rest of the network through the nucleus only and thus are usually not used in communication between the bulk of the network. Inspecting the nodes in this component, we find that they are either leaves (roughly one-third of the nodes) or single nodes that connect with two or more links directly into the nucleus (60%). The rest (8%) form small clusters, with no more than 10 members. Schematically we can imagine these nodes as forming tendrils hanging from the center of the network, as do the tendrils of a jellyfish (see Fig. 2). Because of its Mediterranean origin, we call our decomposition a Medusa model.

Discussion

Although our model is apparently similar to the jellyfish model proposed in ref. 23, our construction and the nature of its parts are different in every detail, as follows. In ref. 23, the nucleus is defined as a maximal clique discovered heuristically (as discussed above), which is found to include only 20–25 ASes. The layers of the jellyfish mantle are nodes labeled by number of hops from the nucleus. Leaves (nodes with a single link) are considered tendrils. Defining any leaf as a tendril does not capture the fact that a significant fraction of the network ($\approx 30\%$) forms small clusters that are in a sense separated from the majority of the network, and that many leaves connect to the peer-connected part of the crust (our “mantle”). Moreover, because the Internet AS graph has a small diameter, the number of layers in the jellyfish mantle is very small. The number of hops from the center is a much less sensitive measure than the k -shell index in our Medusa model. Also, hop count obscures the fractal properties of the Internet that we observe. We therefore conclude our model resolves more precisely the topological structure of the Internet.

In the Medusa model, the distributions of hop count distances between nodes are quite different in each of the three components. In the nucleus, nodes are separated by one or two hops. In the isolated component, nodes connect through the nucleus by paths of two or three steps (which corresponds to one step to the nucleus, optional one step within the nucleus, and one step back to the isolated component). Nodes in the Medusa mantle, or peer-connected component, are typically three or four steps distant from one another using shortest paths in the full network.

Our proposed method of network analysis can be applied to other naturally occurring complex networks as well. Once decomposed, a careful examination of the network components, as the one carried out here for the Internet, can give insight into whether or not the network has the Medusa structure that we find: nucleus, peer-connected component, and isolated tendrils. For example, whereas the actor network (25) shows no tendrils and disconnected k -cores, random scale-free networks do show this structure, but with different quantitative details. The precise values of the different parameters can be used to differentiate models for the Internet.

In ref. 5 it was shown that realistic networks can be divided into two main groups: ones that possess fractal properties, and ones that do not. We show here that the Internet at the AS level, even if initially recognized as a nonfractal network (5), is entirely fractal at the point when it percolates at $k = 6$. Going to higher crusts (i.e., adding more shells), the network is fractal only up to shorter and shorter length scales, until in the limit of the full network, adding up the nucleus, it is completely nonfractal. Therefore, our analysis sheds light on the mechanisms that distinguish fractal and nonfractal networks.

Methods: Distributed Mapping of the Internet

Our Internet topology data sets are among the results of DIMES, a large-scale, distributed measurement effort to measure and track the evolution of the Internet, overcoming the “law of diminishing returns” encountered when measuring the Internet from too few observation points (26). DIMES collects 3–6 million measurements daily from a global network of $>10,000$ software clients. The measurement tool is a lightweight software client, downloaded by $>5,000$ volunteers from the DIMES web site (www.netdimes.org). Each client runs in the background and, a few times every minute, searches out the path to a selected destination elsewhere on the Internet. Destinations are assigned by a central server to each agent, usually at random from a set of ≈ 5 million destinations designed to uniformly cover all Internet Protocol address space in use. More detail is given in ref. 19. All data are logged and analyzed in a data pipeline and added to a database for subsequent analysis.

The data were filtered at two stages. Each client runs each traceroute to a single destination several times (2–4 times in early measurements, later 10 times). When the results differ because of route instabilities, they are discarded, because this effect can give rise to false links. Second, we filtered for certain artifacts resulting from misprogrammed routers by only including edges whose first and last observations were made by different agents, i.e., edges that were seen from multiple locations.

The results of DIMES’ measurements can be analyzed to create several types of topologies, from the router level (where each node represents a single router on the Internet) to the AS topology (where each node is an entire subnetwork, managed by a single organization, usually an Internet Service Provider). This work considers the high-level (AS) topology that results in a network containing $\approx 20,000$ nodes and $\approx 70,000$ links. To obtain the most complete AS graph possible, we supplement the DIMES observations by merging them with the edges exposed by border gateway protocol (BGP) speakers (the software used by ASes to route over long distances is known as BGP), and collected for the past several years by the University of Oregon Route Views Project (www.routeviews.org). About one-quarter of the edges traversed in our experiments are not disclosed in the BGP data sets, and conversely, about one-quarter of the edges disclosed by RouteViews have not yet been seen by DIMES’ agents. The results presented in this article are from a set of measurements conducted from March through May 2005.

A recent series of papers (7, 27) made strong claims that traceroute-based studies from a small number of observers to a large number of destinations would be biased, generating power-law distributions in ordinary random graphs. Their authors argued from simulations that the effect would be difficult to remove by increases in the size of the observer set. However, our observer population is more than two orders of magnitude larger than the 25–50 machines used in the best earlier studies of Internet topology. Although our data are not organized to let us study the progression from few observers to our present numbers, the results, both in terms of power laws and from connectivity analysis, have been quite stable from the beginning of the study to the present. During this time our agent population has grown by four times and has come to spread to >90 countries, in commercial as well as government and academic networks, on all continents. Therefore, we believe that small observer population bias is not an issue with our observations.

We thank S. Solomon, A. Shalit, and A. Vespignani for discussions. S.K. thanks the International Computer Science Institute at the University of California, Berkeley for its hospitality when parts of this work were carried out. The DIMES measurements and our analysis are parts of the

EVERGROW European integrated project 1935, funded within the Sixth Framework's Future and Emerging Technologies division. This work was also supported by the Israel Science Foundation, the Israel

Internet Association, the Israeli Center for Complexity Science, and the European New and Emerging Science and Technology/Pathfinder project DYSONET.

1. Leonardi S, ed (2004) *Algorithms and Models for the Web-Graph: Third International Workshop, WAW2004* (Springer, New York).
2. Pastor-Satorras R, Vespignani A (2004) *Evolution and Structure of the Internet* (Cambridge Univ Press, Cambridge, UK).
3. Dorogovtsev SN, Mendes JFF (2003) *Evolution of Networks: From Biological Nets to the Internet and WWW* (Oxford Univ Press, New York).
4. Albert R, Barabási AL (2002) *Rev Mod Phys* 74:47–94.
5. Song C, Havlin S, Makse HA (2005) *Nature* 433:392–395.
6. Serrano MÁ, Boguná M, Díaz-Guilera A (2005) *Phys Rev Lett* 94:038701.
7. Clauset A, Moore C (2005) *Phys Rev Lett* 94:018701.
8. Eriksen KA, Simonsen I, Maslov S, Sneppen K (2003) *Phys Rev Lett* 90:148701.
9. Cohen R, Erez K, ben-Avraham D, Havlin S (2000) *Phys Rev Lett* 85:4626–4628.
10. Faloutsos M, Faloutsos P, Faloutsos C (1999) *Comput Commun Rev* 29:251–262.
11. Doyle JC, Alderson DL, Lun L, Low S, Roughan M, Shalunov S, Tanaka R, Willinger W (2005) *Proc Natl Acad Sci USA* 102:14497–14502.
12. Pittel B, Spencer J, Wormald N (1996) *J Combin Theory* 67:111–151.
13. Alvarez-Hamelin JI, Dall'Asta L, Barrat A, Vespignani A (2007) e-Print archive, <http://arxiv.org/abs/cs.NI/0511007>.
14. Adler J (1991) *Physica A* 171:453–470.
15. Dorogovtsev SN, Goltsev AV, Mendes JFF (2006) *Phys Rev Lett* 96:040601.
16. Goltsev AV, Dorogovtsev SN, Mendes JFF (2006) *Phys Rev E* 73:056101.
17. Gaertler M, Patrignani M (2004) in *Proceedings of the IPS-Inter-Domain Performance and Simulation*, ed Sallai G (MoMe, Salzburg, Austria), pp 13–24.
18. Batagelj V, Zaversnik M (2002) e-Print Archive, <http://arxiv.org/abs/cs.DS/0202039>.
19. Shavitt Y, Shir E (2005) *ACM SIGCOMM Comput Commun Rev* 35:71–74.
20. Bunde A, Havlin S, eds (1994) *Fractals and Disordered Systems* (Springer, New York), 2nd Ed.
21. Stauffer D, Aharony A (1992) *Introduction to Percolation Theory* (Taylor and Francis, London), 2nd Ed.
22. Cohen R, Erez K, ben-Avraham D, Havlin S (2001) *Phys Rev Lett* 86:3682–3685.
23. Siganos G, Tauro SL, Faloutsos M (2006) *J Commun Networks*, 8:339–350.
24. Goldstein ML, Morris SA, Yen GG (2004) *Eur Phys J B* 41:255–258.
25. Barabási AL, Albert R (1999) *Science* 286:509–512.
26. Barford P, Bestavros A, Byers J, Crovella M (2001) *Proceedings of the First ACM SIGCOMM Workshop on Internet Measurement* (Association for Computing Machinery, New York), pp 5–17.
27. Lakhina A, Byers JW, Crovella M, Xie P (2003) *Proc IEEE INFOCOM* 1:332–341.