

The Connectivity and Fault-Tolerance of the Internet Topology

Christopher R. Palmer* Georgos Siganos[†] Michalis Faloutsos[‡]
Christos Faloutsos[§] Phillip B. Gibbons[¶]

Abstract

In this paper, we apply data mining analysis to study the topology of the Internet, thus creating a new processing framework. To the best of our knowledge, this is one of the first studies that focus on the Internet topology at the router level, i.e., each node is a router. The size (280K nodes) and the nature of the graph are such that new analysis methods have to be employed. First, we suggest computationally-expensive metrics to characterize topological properties. Then, we present an efficient approximation algorithm that makes the calculation of these metrics possible. Finally, we demonstrate the initial results of our framework. For example, we show that we can identify “central” routers, and poorly connected or even isolated nodes. We also find that the Internet is surprisingly resilient to random link and router failures, having only small changes in the connectivity for fewer than 10,000 failures. Our framework seems a promising step towards understanding and characterizing the Internet topology and possible other real communication graphs such as web-graphs.

1 Introduction

In this paper, we study the topology of the Internet at the router level. We know very little about the Internet, despite the significance and impact of the network in everyday life. This is especially true for the topology of the network, which is a crucial part of modeling and simulating the network. First, we study the structure of the network. Using topological properties, we manage to identify “different” parts of the network such as central backbone routers and areas with poor connectivity. Second, we study the robustness of the topology to edge and node failures. In our study, we use a novel data mining tool to process the large topological data.

Why can’t we model the Internet topology? There are several reasons for that. First, the necessary data has only recently become available. Second, the data is so large (285K nodes) that standard processing and visualization techniques are inadequate. Despite these challenges, it is very important to characterize the topology. The absence of this knowledge is one of the reasons “Why we don’t know how to simulate the Internet” according to Paxson and Floyd [12]. It is very

*Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213

[†]U.C. Riverside, Dept. of Comp. Science Riverside, CA 92521

[‡]U.C. Riverside, Dept. of Comp. Science Riverside, CA 92521

[§]Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213

[¶]Information Sciences Research Center, Bell Laboratories, Murray Hill, NJ

difficult to analyze and optimize network performance without understanding its topology. It is analogous to attempting to resolve traffic problems in a city without a map.

The novelty of our work lies in the use of data mining tools in the study of the network topology. In more detail, we use a new tool to approximate the size of the neighbourhood of a node. This is a computationally expensive procedure and it is made feasible by our tool that reduces its runtime by more than a factor of 400. Given this tool, we are able to provide results in two main directions. First, we provide new insight in the structure of the graph, and we classify nodes according to their neighbourhood related properties. For example, we show that using our analysis we can distinguish nodes which may correspond to centrally-located or backbone routers. Second, we study the robustness of the network to edge and node failures. We find that the network is robust to edge failures, and uniformly distributed node failures. However, we observe that failures in the central or backbone nodes can very quickly hurt the connectivity of the network. In our work, we use some of the novel graph metrics proposed by Faloutsos et al [4] and we show that these metrics can actually provide insight into the graph structure.

In section 2 we present the background, summarize previous work and define our novel Internet metrics. In section 3, we highlight the approximation algorithm that we use. In section 4, we present our results. Finally in section 5, we present our conclusions.

2 Background and Novel Graph Metrics

We begin this section by defining the terminology that we will use in the remainder of the paper. We also describe the Internet router data that we will be using.

We study the network at the *router level*, that is, each Internet router is represented by a node in the graph while each link is mapped by an edge. In contrast, a lot of previous work has concentrated on the interdomain level of the topology where each graph node represents a domain or Autonomous System. The router level is a much larger and more detailed graph. We believe that the size of this graph has made its processing prohibitively expensive.

We list a number of graph metrics that have been proposed in the literature only recently. Typical metrics are average node degree and diameter. We believe that the *hop exponent*, *effective diameter* and *effective eccentricity* are much more effective in characterizing the complexity of the Internet graph.

Definitions. Let $G = (V, E)$ be either a directed or undirected graph. Let $d(u, v)$ be the shortest path distance from node u to node v . We define the following :

Reachable set: Nodes that are within distance h of u : $S(u, h) = \{v : d(u, v) \leq h\}$.

Individual neighbourhood function: Reachable sets sizes: $N(u, h) = |S(u, h)|$.

Neighbourhood function: Number of pairs of nodes within dist. h : $N(h) = \sum_{u \in V} N(u, h)$.

Reachable pairs: Number of pairs of nodes that have a path connecting them: $N(\infty)$.

Effective Diameter: Least distance, h , such that at least 90% of the reachable pairs are

within distance h : $\min_h N(h) \geq .9 \cdot N(\infty)$

Effective Eccentricity: Eccentricity is to a node as the diameter is to a graph. Least distance, h , such that 90% of u 's reachable set are within distance h : $\min_h N(u, h) \geq .9 \cdot N(u, \infty)$.

Hop-Plot Exponent: A proposed power-law in [4], that the total number of pairs of nodes within h hops is proportional to the number of hops raised to a constant, \mathcal{H} (*hop-plot exponent*).

We use the hop-plot exponent to characterize the growth of the neighbourhood function. To compute it, we apply log transforms to the neighbourhood function and compute the least-square fitting line for the points up to the effective diameter.

Our real Internet graph. We use the graph that is the result of the union of the SCAN [15] and Lucent Internet mapping project results [14]. The SCAN project develop a topology discovery tool called Mercator that uses hop-limited probes – the same primitive used in traceroute – to infer the map of the Internet [15]. The Lucent Internet mapping project uses a single probe location but has performed long term monitoring [8]. This merged data set represents the best map of the Internet (at a router level) which was current as of late 1999. The resulting graph has approximately 285K nodes, 430K edges, a maximum degree of 1,978 and an average degree of 3.15. It is this graph that we will use to study the router-level Internet topology.

Previous Work. There have been several measurements of the Internet topology [6, 11, 7]. These studies focus on the collection of data while the analysis appears secondary. There has not been yet a comprehensive study of the Internet topology at the router level. In contrast, the interdomain level has been studied lately [16]. In a parallel tangent, several people have studied topological properties indirectly, through the study of scaling of multicast trees in Internet [2, 13, 17]. Recently, Albert et al. [1] and Tauro and Faloutsos [16] studied the fault tolerance of the Internet at the interdomain level. Recall that our work here focuses at the router level of the Internet.

Assumptions and Limitations. Our work relies on measured data which is always subject to measuring errors. The main problems with Internet measurements are a) incompleteness, b) router identification. We can not claim or guarantee that it is most of it, but we have reasons to believe that this data contains a substantial and representative part of the Internet. For a discussion on this issue see [15, 7, 8].

3 Approximate Neighbourhood Function (ANF)

We have developed and evaluated an approximate neighbourhood function (ANF) in [10]. We present the key ideas and a simple version of the algorithm. The algorithm is presented in sufficient detail to reproduce the results in this paper. The underlying approach is to iteratively compute $S(u, h)$, the set of nodes within at most h hops of u , using the edge set and $S(u, h - 1)$. That is:

```
FOR each node u DO S(u,0) = { u }
FOR each iteration, h starting at 1
  FOR each node u DO S(u,h) = S(u,h-1)
  FOR each edge (u,v) DO S(v,h) = S(v,h) U S(u,h-1)
```

```

FOR each node, u DO
  M(u,0) = concatenation of k bitmasks, each with 1 bit set
           (according to an exponential distribution)

FOR each distance, it, starting with 1 DO
  FOR each node, u DO M(u,it) = M(u,it-1)
  FOR each edge (u,v) DO M(u,it) = (M(u,it) OR M(v,it-1))

The estimate is: SUM(all u) (2^b)/(.7731*bias)
  where b is the average of the least zero bits in the k bitmasks
  bias, a small bias factor, is (1+.31/k)

```

Figure 1: In-Core Approximate Neighbourhood Function (ANF)

and then the neighbourhood function is $N(h) = \sum_u |S(u,h)|$. This algorithm will be horribly inefficient to use in practice because the set operations are expensive. Instead, we use a tool called approximate counting. An *approximate counting* algorithm takes as input a multi-set and then estimates the number of distinct elements in the multi-set. In [5], each possible element (for us, that is each node) is assigned a random bit using an exponential distribution (half the nodes get bit 0, a quarter get bit 1, etc). To estimate the number of elements in a multi-set, you simply OR together the bits that we assigned to each element. The estimate is then close to 2^b , where b is the least zero bit in the bitmask. We can use the approximate counting idea to replace the set operations in our simple algorithm. We use $M(u,h)$ to denote the bitmask approximation to $|S(u,h)|$ and to improve accuracy perform the approximation k times in parallel (we fix $k = 64$ for the remainder of the paper). This algorithm appears in Figure 1 and operates only in-core. An external version of the algorithm is also presented in [10] which allows the processing of graphs that are much larger than available memory.

The ANF algorithm has several properties that make it possible to perform studies on large graphs. Here we will just list its properties, while appendix A provides a sense of its efficiency highlighting some of the actual executions of the algorithm [10].

- **Fast:** Approximate the neighbourhood function of the router level graph (285K nodes and 430K edges) in a matter of minutes, rather than nearly a day (436x faster).
- **Accurate:** The approximates have provable bounds, generally within 5-10% of the true function for our experiments.
- **Individual neighbourhood functions:** As a by-product of computing the neighbourhood functions, we compute the neighbourhood function for each individual node.

4 Data-Mining on the Internet Graph

In this section, we show how the new approximation algorithm enables us to find interesting properties of the Internet topology. First, we study the structure of the network by identifying properties of the nodes. Second, we study the robustness of the topology to component failures.

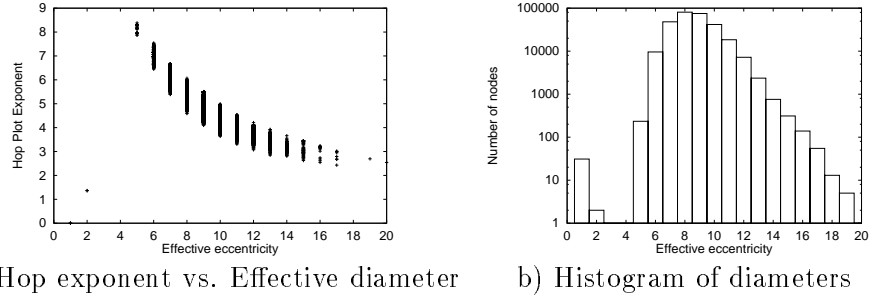


Figure 2: Hop exponent and effective eccentricities of the individual neighbourhood functions for Internet routers

4.1 Discovering structure through node classification

The first application of the neighbourhood function is to examine the individual routers in the Internet graph. We wish to understand the connectivity richness of different routers. Note that this corresponds loosely to alternate paths that the node could potentially use. That is, we know that the effective diameter of the Internet is approximately 10 hops, but how does that translate to the path lengths for individual nodes.

Quantifying the relationship of hopplot and eccentricity. We conducted an experiment with results in Figure 2 a). First we compute the individual neighbourhood functions for all nodes in the graph. From this, we measure the effective eccentricity of each node and compute the hop exponent for each individual neighbourhood function. The scatter plot of these values indicates a strong correlation between the eccentricity and the hop exponent (not surprising).

Classifying nodes using eccentricity. In Figure 2 b), we show the histogram of the number of nodes with each effective eccentricity. The number of nodes is plotted in log scale. We observe that about 10,000 routers have an effective eccentricity of at most 6 and another 10,000 routers have an effective eccentricity of 12 or larger. The majority of the nodes have an effective eccentricity that is close to the effective diameter of the Internet.

Identifying pathologies of the measured data. Our analysis can help identify pathological or incomplete cases in the measured data. We observe that there are some nodes with eccentricity 1 and 2. This would mean that a node has a degree of approximately $0.9 \cdot 284K = 250K$ links. This is clearly not the case, since the max degree of the graph is no more than 2000. The explanation is that the graph has some disconnected components. Recall that we define the eccentricity relative to the nodes that the *can be* reached. For example, an isolated pair of nodes has eccentricity 1. We checked the data in the graph and found that this was actually the case for the nodes reported here.

4.2 Topological Fault-tolerance

When looking at the failure behaviour of a network, there are two contributing factors. First, network outages may disconnect pairs of nodes. Second, protocol failures may cause pairs of nodes

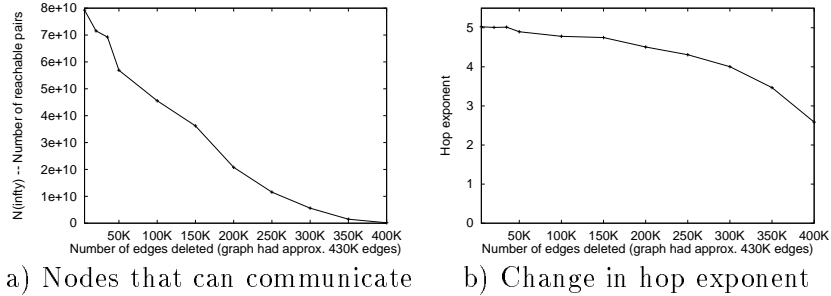


Figure 3: Effect of edge deletions (connection failures) on the router graph

to appear disconnected even though a physical path exists between them. In this section, we explore the first form of failure. Network outages take the form of either connections being lost (for example, as a result of a back-hoe breaking a network cable) or a computer failing.

4.2.1 Link Failures

We want to examine the robustness of the Internet with respect to link failures. Thus, we conduct the following experiment. We select x edges at random and delete them. Then, we compute the neighbourhood function and measure the effective diameter, the number of reachable pairs, and the hop exponent.

The Internet is robust to link failures. The number of reachable pairs, and the hop exponent shown as a function of the number of edge deletions in Figure 3. Each point is the average of over 3 randomly chosen sets of edges. We have not shown the average diameter as it only varies from 10 to approximately 12 over the full set of edge deletions. This shows that while we delete edges, pairs of nodes either become disconnected or they have an alternative path that is not significantly longer. We see that the Internet is quite resilient under connection failures, with only a small decrease in the number of reachable pairs and the hop exponent for fewer than 50,000 failures. Moreover, deleting edges does not appear to change the hop exponent until closer to 200,000 deletions, suggesting that while we partition the Internet, the structure is preserved.

4.2.2 Node Failures

Router failures represent a more catastrophic event, since a router affects all its adjacent edges. Therefore, we expect that this will cause problems at least in the vicinity of the node. To model a router failure, we select a node (in one of three different ways) and delete all its adjacent edges.

We introduce failures in three different ways. First, we randomly select routers (uniform distribution). This corresponds to an unbiased router-specific failure. Second, we take the opposite of the previous approach. We remove the nodes in order of highest degree. This is the most aggressive approach that we can take to decompose the Internet. The final method that we use removes nodes in order of highest individual hop exponent. We theorized in section 4.1 that the nodes with low

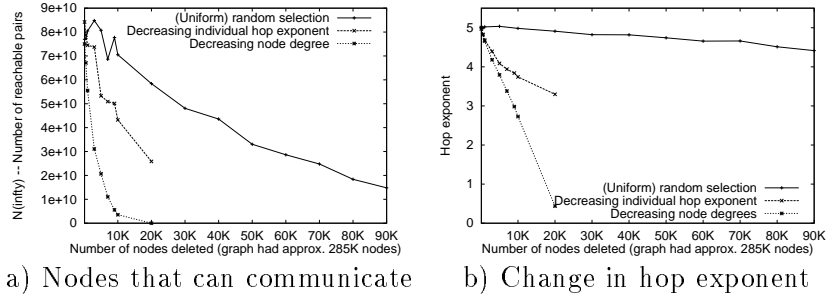


Figure 4: Effect of node deletions (router failures) on the router graph

effective diameter and high hop exponent may be important nodes, possibly in the backbone. These results are shown in Figure 4.

The Internet is robust to random node failures. Here we see that fewer than about 10,000 router failures does not significantly affect the Internet. That is, the number of reachable pairs is not drastically reduced. Moreover, the hop exponent is not decreasing very quickly with the number of deleted nodes. This suggests that it might be possible to use sampled Internet graphs to conduct reliable simulation experiments.

The Internet is sensitive to focused failures. The other two approaches for router selection have much more devastating results. By selecting routers according to their degree, we can quickly disconnects the Internet. Only 10,000 nodes are needed to effectively disconnect most node pairs. Even 100 of these nodes is sufficient to remove the connectivity in more than 5 billion pairs.

The effective eccentricity as a node classifier. The above observation confirms that the effective eccentricity is a useful graph metric. Nodes with high eccentricity are “important” for the network; their failures create problems to the connectivity of the network. At the same time, we see that eccentricity has a different effect than the degree regarding the connectivity. This observation suggests that eccentricity quantifies another aspect of the “importance” of the node and it is distinct from the node degree.

5 Conclusions

In this paper we proposed a set of new and existing metrics that can capture interesting topological properties of the Internet. While these metrics are expensive to compute exactly, we have shown that a new data mining tool can be very effective in approximating these metrics. This has allowed us to obtain new insights into the properties of routers and their “importance” for the network. Further use of these ideas allowed us to examine the intrinsic resilience of the Internet.

Our work highlights the untapped potential that exists for the use of data-mining tools in network data; the size of the network makes classic techniques prohibitively expensive in computation time. In more detail our results can be summarized in the following points:

- We propose a new set of new and existing metrics capturing interesting topological properties.

- We show that a new data mining tool can be very effective in calculating the otherwise computationally expensive metrics.
- Effective eccentricity is a good node metric, for of the “topological significance” of a node and it actually captures a different aspect than that of the node degree.
- The Internet topology is resilient to random link and node failures.
- The Internet is sensitive to focussed node failures; it is sensitive to failures of nodes of high “significance” expressed either by its degree or effective eccentricity.

Our metrics and tool make a promising step towards understanding and characterizing the Internet topology and possible other real communication graphs such as web-graphs. We are in the process of developing more methods to decipher the structure of the Internet topology. Our initial results are a promising step.

References

- [1] R. Albert, H. Jeong, and A. Barabasi. Attack and error tolerance of complex networks. *Nature*, 406, July 2000.
- [2] J. Chuang and M. Sirbu. Pricing multicast communications: A cost based approach. *In Proc. of the INET'98*, 1998.
- [3] CORA search engine. <http://www.cora.whizbang.com>.
- [4] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. *In SIGCOMM*, 1999.
- [5] P. Flajolet and G. N. Martin. Probabilistic counting algorithms for data base applications. *Journal of Computer and System Sciences*, 31:182–209, 1985.
- [6] R. Govindan and A. Reddy. An analysis of internet inter-domain topology and route stability. *Proc. IEEE INFOCOM*, Kobe, Japan, April 7-11 1997.
- [7] R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery. *Proc. IEEE INFOCOM*, Tel Aviv, Israel, March 2000.
- [8] <http://cm.bell-labs.com/who/ches/map/index.html>.
- [9] C. R. Palmer and J. G. Steffan. Generating network topologies that obey power laws. *In IEEE Globecom 2000*, 2000.
- [10] Christopher R. Palmer, Phillip B. Gibbons, and Christos Faloutsos. A fast approximation of the “neighbourhood” function for massive graphs. Under review.
- [11] J.-J. Pansiot and D Grad. On routes and multicast trees in the Internet. *ACM Computer Communication Review*, 28(1):41–50, January 1998.
- [12] V. Paxson and S. Floyd. Why we don’t know how to simulate the internet. *Proceedings of the 1997 Winter Simulation Conference*, December 1997.
- [13] G. Philips, S. Shenker, and H. Tangmunarunkit. Scaling of multicast trees: Comments on the chuang-sirbu scaling law. *ACM SIGCOMM*, Sep 1999.
- [14] <http://www.isi.edu/scan/mercator/maps.html>.
- [15] http://www.isi.edu/scan/Pubs/scan_proposal.ps.gz.
- [16] S.L. Tauro and M. Faloutsos. Fault-tolerance and robustness of the internet topology. *PRDC 2000, (abstract)*, Los Angeles, 2000.
- [17] T. Wong and R. Katz. An analysis of multicast forwarding state scalability. *International Conference on Network Protocols*, 2000.

Table 1: Example graphs and their salient parameters

Graph	# Nodes	# Edges	Eff. Diameter	Directed?	Running Time (min)		
					Approximation	Exact	Speed-up
Cora	127,083	330,198	28	Yes	1.5	6	4x
80-20	166,946	449,832	8	No	1.5	680	453x
Router	284,805	430,342	10	No	2.75	1,200	436x




Figure 5: Actual neighbourhood function and our excellent approximations. Also shown are the least-squares line used for generating the hop exponent.

A The efficiency of our approximation algorithm

We highlight the accuracy and speed of the approximation algorithm. Note that a thorough analysis is beyond the scope of this paper. For such an analysis see [10]. We show three examples to anecdotally illustrate the quality of approximation that we are expecting. We take three data sets that have different sizes and properties (see Table 1). First, the Cora data set is a directed graph extracted from a large collection of computer science papers (each node is a paper and each directed edge is a citation) [3]. Second, the 80-20 generator creates Internet-like topologies [9]. Finally, the Router data set was explained earlier. For each data set, we compute the exact neighbourhood function using repeated breadth-first searches (which is state of the art), compute our approximation and compute the least squares fitting line based on our approximation. These are plotted in log – log scale in Figure 5. We see that all approximations are excellent and that the line provides an excellent characterization of the growth of the neighbourhood function. This approach appears reasonable for both directed and undirected graphs and for very different diameters. Finally, we see that ANF runs more than 400 times faster than the full computation for our Internet type graphs.