

# Bounding Damage From Link Destruction, with Application to the Internet

George Dean Bissias, Brian Neil Levine, and Arnold Rosenberg  
Dept. of Computer Science, University of Massachusetts Amherst  
Amherst, MA 01003, USA  
{gbiss,brian,rsnbrg}@cs.umass.edu

## Categories and Subject Descriptors

C.4 [Performance of Systems]: Reliability, availability, and serviceability; C.2.0 [General]: Security and protection

## General Terms

Security, Reliability, Human Factors

## Keywords

Spectral Graph Theory, Graph Partitioning, Vulnerability

## 1. INTRODUCTION

If an attacker could remove any  $k$  links from a network, how would connectivity suffer in terms of the number of pairs of communicating nodes? The answer to this question is a critical component of quantifying the damage to networks that results from attacks on inter-domain routing, intra-domain routing, mobile network routing, and sensor networks. In this short paper, we introduce a technique that provides a lower bound on a graph's connectivity after link removal.

Our experimental results show that tight bounds can be expected for attacks involving the removal of a small fraction of the network links. While this is a limitation, we expect that, in general, the resources required to carry out an attack scale proportionally with the number of links removed. Consequently, our results are tight for cases that we expect to be more common since weaker attackers are more numerous.

## 2. SPECTRAL LOWER BOUND

In this section, we sketch an algorithm that finds the *resilience* of a graph. The resilience bounds a graph's connectivity from below after the targeted removal of a fixed number of edges. Connectivity is defined as the sum of squares of the sizes of the connected components  $\mathcal{C}(G)$  of  $G$ . Intuitively, if  $G$  represents a communication network, then the connected components represent the regions of the network that are in communication. The resilience of a communication network is the lowest number of communicating entities that will exist after removing a fixed number of links in the network. The following theorem due to Donath and Hoffman [3] serves as a point of departure for our bound.

This work was supported in part by the National Science Foundation under grants CNS-0133055 and ANI-032586.

Copyright is held by the author/owner(s).  
SIGMETRICS'07, June 12–16, 2007, San Diego, California, USA.  
ACM 978-1-59593-639-4/07/0006.

**THEOREM 1.** *Let graph  $G$  be given. For an arbitrary  $\Omega \in E$  and arbitrary ordering,  $\{C_1, \dots, C_k\} = \mathcal{C}(G^\Omega)$ , of the connected components of  $G^\Omega$ :  $|\Omega| \geq \frac{1}{2} \sum_{i=1}^k m_i \lambda_i = \frac{1}{2} \mathbf{m}^T \boldsymbol{\lambda}$ , where  $G^\Omega$  is  $G$  with edges  $\Omega$  removed,  $\boldsymbol{\lambda}$  is comprised of the  $k$  smallest eigenvalues of the Laplacian of  $G$ , and  $m_i = |\mathcal{V}(C_i)|$  (i.e., the size of  $C_i$ ).*

From this we derive an optimization problem.

**PROBLEM 1.** *For fixed  $\omega \in \mathbb{N}$  and variables  $y \in \mathbb{R}$ ,  $k \in \{1, \dots, \omega + 1\}$ , and  $\mathbf{m} \in \mathbb{N}^k$  define  $f : \mathbb{N}^k \rightarrow \mathbb{R}$ ,  $f(\mathbf{m}) = \mathbf{m}^T \mathbf{m}$ ;  $g : \mathbb{N}^k \times \mathbb{R} \rightarrow \mathbb{R}$ ,  $g(\mathbf{m}; y) = \frac{1}{2} \mathbf{m}^T \boldsymbol{\lambda} - y^2 - \omega$ ;  $h : \mathbb{N}^k \rightarrow \mathbb{R}$ ,  $h(\mathbf{m}) = \mathbf{m}^T \mathbf{1} - n$ . With  $n$  equal to the number of vertices in  $G$ ,  $k$  corresponding to the maximum number of components, and  $y$  being slack, minimize  $f(\mathbf{m})$  subject to constraints  $g(\mathbf{m}; y) = h(\mathbf{m}) = 0$ .*

The solution to the continuous relaxation of this problem is also a solution to the resilience problem. It can be shown that a solution can be found in time cubic in the number of vertices in  $G$ .

## 3. RESULTS

In this section, we demonstrate the quality of our bounds by applying them to a measured Internet graph and several synthetic graphs often compared against Internet data because of their salient characteristics.

**3.1 Graphs** The CAIDA skitter project maps Internet autonomous systems (ASes) by using traceroutes to reconstruct the AS level topology. We examined 36 recent skitter graphs, recorded on the 1st, 10th, and 20th of each of the months between January 2006 and December 2006, inclusive. Each graph comprised tens of thousands of nodes denoting ASes. The graphs were also highly disconnected, so we chose to focus on the largest connected component of each. This largest subgraph typically had approximately 16,000 links and 8,000 nodes for each graph. For convenience, we refer to the set of these 36 graphs as the CAIDA graphs. We also generated three sets of synthetic graphs with roughly the same numbers of nodes and links as the CAIDA graphs. 36 Power-law graphs, labeled BA, according to the preferential attachment model of Barabasi and Albert [1], 36 HSF graphs according to the procedure outlined in Li et al. [5], and 36 HOT graphs, also defined in Li et al. [5].

**3.2 Experiments** For each of the graphs in CAIDA, BA, HSF, and HOT, we generated an upper bound on the

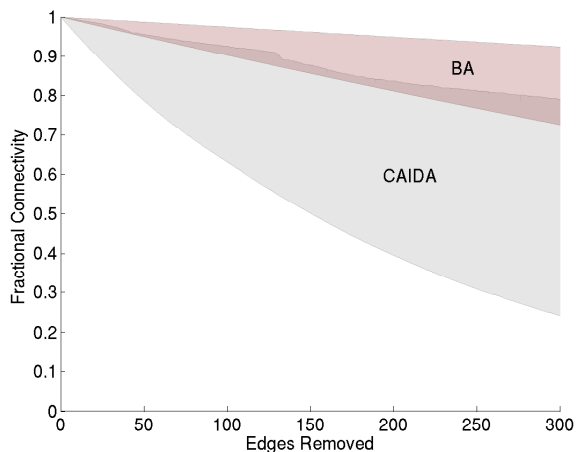


Figure 1: Median value Upper and Lower bounds for the CAIDA and BA graphs

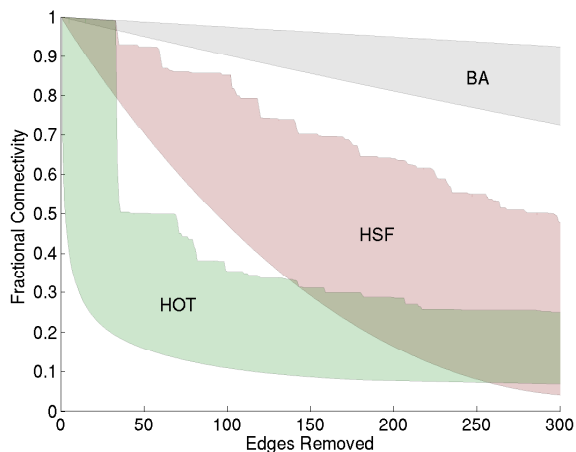


Figure 2: Median value Upper and Lower bounds for the BA, HSF, and HOT graphs

their resilience using the Metis graph partitioning library, and a lower bound using the technique outlined in Section 2. With the exception of HOT, the upper and lower bounds for each graph were closer for smaller numbers of links removed. Nevertheless, CAIDA, HSF, and BA stayed within a factor of three for the first 260 links removed with BA remaining well within a factor of two. Figures 1 and 2 show the median upper and lower bound values for each of the graph types.

The bounds for CAIDA, as shown in Figure 1, weakened somewhat rapidly but they did generally indicate that communication was easy to disrupt and difficult to completely destroy. The upper bound represents demonstrable damage from removing a specific set of links, and it shows, for instance, that the removal of just 25 specific links out of roughly 16,000 results in a loss of about 1.75% of all AS-to-AS connectivity. On the other hand, the lower bound states that the removal of no combination of 25 links will reduce connectivity more than 12%. From another perspective, after 150 links are removed from CAIDA, about 50% of ASes can still communicate.

HOT displays a unique decay in connectivity in the upper bound, as show in Figure 2. After approximately 40 links are removed, apparently strong robustness yields abruptly, resulting in low connectivity that closely matches the lower bound. This is similar to what was observed in Li et al. [5] where it was found that the most effective node attack against a similarly constructed HOT graph was to remove core nodes.

**3.3 Comparing Graph Resilience** Having the same number of nodes and roughly the same number of links, it makes sense to examine the relative resilience of each graph. In Figures 1 and 2, BA demonstrates higher or equal resilience compared to all other graphs. That BA is at least as resilient to link attack as CAIDA (for small numbers of links removed) has not yet been reported to the best of our knowledge. The figures also show that CAIDA is much more resilient than HOT for 50 to at least 260 links removed. This is attributable to the fact that HOT graphs were developed to model specific network structures [5, 4], but the implication here is that while a simple HOT graph may exhibit alarming vulnerability, a more complex HOT graph, CAIDA (by definition not construction), does not.

The contrast between the three network models, BA, HOT, and HSF is also interesting. Figure 2 shows bounds for each graph type overlaid together. For removed link quantities ranging from approximately 40 to 140, HSF clearly exhibits higher resilience than HOT. BA exhibits higher resilience than either HOT or HSF from approximately 50 links removed to at least 260. It's possible that this illustrates the effect of *inadvertent weakness* first noted by Carlson and Doyle [2] whereby structural weaknesses are inadvertently introduced in the process of engineering a network.

## 4. CONCLUSION

We have introduced a technique to bound graph connectivity from below after targeted link removal. We applied our bound to graphs created from CAIDA measurements of the Internet AS-level topology as well as BA, HOT, and HSF synthetic graphs. In our experiments, we show several quantitative comparisons between the four types of graphs. Our results model the most common attacker with few resources. We found that the CAIDA AS level graph is no more resilient than the BA graph, but after removal of 150 links, the majority of AS pairs can still communicate for any link attack. The HOT graph is shown to be distinctly more vulnerable than the CAIDA, HSF, and BA graphs.

## 5. REFERENCES

- [1] BARABASI, A.-L., AND ALBERT, R. Emergence of scaling in random networks. *Science* 286 (1999), 590–512.
- [2] CARLSON, J. M., AND DOYLE, J. Highly optimized tolerance: a mechanism for power laws in designed systems. *Physics Review E* (1999).
- [3] DONATH, W., AND HOFFMAN, A. J. Lower bounds for the partitioning of graphs. *IBM Journal of Research and Development* 17 (1973), 420–425.
- [4] DOYLE, J. C., ALDERSON, D. L., LI, L., LOW, S., ROUGHAN, M., SHALUNOV, S., TANAKA, R., AND WILLINGER, W. The "robust yet fragile" nature of the Internet. *Proceedings of the National Academy of Sciences of the United States of America* (2005).
- [5] LI, L., ALDERSON, D., TANAKA, R., DOYLE, J. C., AND WILLINGER, W. Towards a theory of scale-free graphs: definition, properties, and implications (extended version). Tech. rep., California Institute of Technology, 2005.