

Network Heterogeneity and Cascading Failures – An Evaluation for the case of BGP Vulnerability

Christian Doerr¹, Paul Smith² and David Hutchison²

¹ Department of Telecommunication, TU Delft, The Netherlands
c.doerr@tudelft.nl

² Computing Department, Lancaster University, UK
{p.smith,dh}@comp.lancs.ac.uk

Abstract. Large-scale outages of computer networks, particularly the Internet, can have a significant impact on their users and society in general. There have been a number of theoretical studies of complex network structures that suggest that heterogeneous networks, in terms of node connectivity and load, are more vulnerable to cascading failures than those which are more homogeneous. In this paper, we describe early research into an investigation of whether this thesis holds true for vulnerabilities in the Internet’s inter-domain routing protocol – BGP – in light of different network structures. Specifically, we are investigating the effects of BGP routers creating blackholes – observed phenomena in the Internet in recent years. We describe our evaluation setup, which includes a bespoke topology generator that can fluidly create any topology configuration from the current scale-free AS-level to the investigated homogeneous graphs. We find that network homogeneity as suggested by theory does not protect the overall network from failures in practice, but instead may even be harmful to network operations.

1 Introduction

Recently, a multitude of studies have been conducted that investigate under which conditions complex network structures are most vulnerable. These studies conclude that scale-free networks are relatively robust against random failure and intentional attacks [1, 2]. The failure or breakdown of a network component and its effects may not remain self-contained, but depending on the network characteristics and model, might spread across the infrastructure, resulting in cascading failures in a large part of the system [3]. Many of the studies investigating cascading failures suggest that high levels of heterogeneity, i.e., large differences in node connectivity and traffic share among the network nodes, will increase the vulnerability and worsen the outcome of a failure cascade, and that structural homogeneity will dampen the effects.

While collapses of the Internet (or parts thereof) due to traffic overload, as described in [3], have not been observed in reality, cascading failures and subsequent inoperability of some parts of the Internet have frequently been documented for the case of BGP, which is vulnerable to unintentional misconfiguration and could theoretically also be exploited in intentional attacks [4].

In this paper, we describe early research that aims to merge the theoretical findings on abstract scale-free networks with AS-level structural and BGP behavioral data, and investigate whether structural network heterogeneity is indeed a risk for cascading problems, and may be alleviated when the network would be designed more homogeneous.

2 BGP Vulnerabilities

We will focus on two types of BGP vulnerabilities that have been readily observed in the Internet: blackholing via false prefix announcements and interdomain routing instabilities via route flapping. Before we describe these vulnerabilities, we sketch the general functioning of the BGP protocol, and how it enables interdomain routing. For an in-depth discussion of BGP, we refer the reader to [5].

In simple terms, what is most commonly referred to as “the Internet” is a conglomeration of individual networks which have been *interlinked* into a larger network (see Fig. 1). Individual networks are referred to as Autonomous Systems (ASes), since they are maintained by a single administrative entity, such as an Internet Service Provider (ISP). Currently, approximately 27,000 ASes provide global connectivity.

Before this can happen, it is necessary to compute paths along which remote parts of the network may be reached. This is the task of the BGP protocol, through which each AS shares information describing its own connectivity with neighbors and announces currently available routes through its own domain to others. In the example of Fig. 1, AS1 may announce to its neighbors that it has direct connections (1 hop away) to AS 3, 4 and 5. As the other neighboring systems do the same, AS1 can further learn that the IP ranges of AS4 may also be reached through AS3, but through a longer path of length 2, and that AS3 will also provide connectivity to AS2, which in turn can forward data to AS6. While BGP can be configured to prefer specific routes and neighbors over others, it typically chooses the shortest and most specific³ route to its destination. Routes can be advertised or withdrawn as the network topology changes.

³ the smallest and thereby most specific description of an IP block




Fig. 1. An AS topology vulnerable to blackholing and route flapping.

Blackholing While the mechanism of BGP advertisements and withdrawals is designed to be reliable, problems can occur when routers announce wrong routing and reachability information due to accidental or intentional misconfiguration. Consider the situation where AS1 falsely advertises that it has a (new) direct route of length 1 to the IP range of AS 6. This announcement would let AS1's neighbors, AS5 and AS4, send any traffic destined for AS6 to routers D and E in AS1, as this new route of length 2 is shorter than the previously known path of length 4 and 3, respectively. In the area where AS1's announced path is superior to the correct advertisement of AS2, a blackhole is created that disconnects AS6 from this part of the network.

Such situations, where networks falsely announce IP prefixes that they do not own thus leading to blackholes, happen quite frequently in the Internet. In an analysis of historical data, Zhao et al. [6] showed that between 1998 and 2001 about 30 conflicting routes occurred daily. While most of these issues go unnoticed for the vast majority of the Internet, there exist several instances where false advertisements had a global impact on the network and disrupted connectivity on a planetary scale: in 1997, a false advertisement of AS7007 [7] resulted in that network being recognized as the best path to almost the entire Internet. Similar events on a smaller scale were repeated in 1998 and 2001. The most well known recent blackhole event may be the "YouTube incident", where an intentional BGP misconfiguration by Pakistan Telecom to block YouTube in Pakistan was leaked outside its network and disconnected the website from large parts of the Internet [8].

Route flapping Even when each AS only propagates correct routing information, another type of routing problems can occur. Due to the long convergence times of BGP, routers can be configured to employ a mechanism called route dampening to protect the network from routing instabilities. When remote nodes repeatedly announce and withdraw routes in short succession, a large configuration overhead of finding and selecting new routes is introduced throughout the network. Routers will therefore temporarily remove paths from the routing tables when select routes begin to flap, resulting in these nodes becoming unreachable until the dampening expires.

While this mechanism is intended to improve the convergence of BGP and its routes when facing unstable paths, this behavior can be exploited [9]. If router G of AS54 in Fig. 1 would deliberately make repeated path announcements and withdrawals which are then propagated by its neighbor AS1, other upstream networks as such AS2, 3 and 4 may enable route dampening to protect the network, temporarily block these flapping routes, and thereby disconnect AS1 from the network.

Both these error types can create large-scale cascading effects as remote nodes react to the incoming route updates. As discussed earlier, the theoretical analysis of cascading failures in complex networks has resulted in the finding that cascades may be prevented, or at least limited in scope, as the network topology structure becomes more homogeneous. Next, we discuss how we propose to test this claim for the Internet as a concrete complex network.




Fig. 2. Degree-sorted circular layouts visualize the characteristic degree distributions of random and scale-free graphs.

3 Evaluation Setup

In order to evaluate whether number and magnitude of cascading failures within the BGP protocol could be reduced, if not avoided, by higher levels of topological homogeneity, we implemented a testing suite that could create, emulate and test a large number of networking scenarios. This testing suite was based upon SSF-Net [10], an open-source discrete event simulator for complex systems with nanosecond-resolution emulation capabilities of all the protocols necessary for this evaluation. This simulator was extended with custom-made modules to mimic the behavior of a malicious agent conducting attacks against random parts of the network.

The simulator was additionally connected to a custom-made topology generator which was used to create customized Internet topologies. Since the goal of our study was to compare current Internet-like topologies with potential alternative, more homogeneous AS topologies, the topology generator was developed in such a way that it could fluidly create any type and mixture of network scenario between these two opposite poles, a scale-free AS-level or a purely random graph. While the procedures of creating good random Erdős-Renyi graphs has been sufficiently explored, there still exists much debate over the best algorithms to create a network topology that can mimic the topology of the Internet. For our evaluation, we implemented a scale-free topology generator from the PLRG class, which has been found to generate network structures with characteristics closely matching those of the Internet at the AS-level [11]. Since the underlying generation procedure of the models is identical, it is possible to overlay and combine the two models into a single generator, where a single parameter ($s \in [0, 1]$) then specifies the level of dominance of one model over the other, where $s=0$ indicates the full dominance of the random graph and $s=1$ the full dominance of the scale-free topology. Fig. 2 shows three example topologies with 50 nodes and 200 links for the three parameter values $s=0, 0.5, 1$, arranged in a degree-sorted circular layout. It is easy to see the powerlaw distribution of node degrees in the case of $s=1$ and the complete absence of patterns (i.e., randomness) for $s=0$. The

hybrid graph with $s=0.5$ represents a mixture between the two models, where 50% of the links are created according to a scale-free topology, while the other 50% are added randomly.

4 Results

Due to space constraints, this section describes only a selection of results obtained from the simulations, and focuses specifically on the blackholing attack.

As discussed before, the main objective of this work was to evaluate the hypothesis obtained from theoretical modeling that network homogeneity will hinder the spread of failures. To test this prediction, we generated topologies with varying sizes $n=100..5000$, where the nodes were interconnected according to the scheme discussed in Sect. 3. For each network size, 7 classes of topologies were created, defined by the s -value $s=0, 0.1, 0.25, 0.5, 0.75, 0.9, 1$. In the SSF-Net simulation environment, each node of the graph was configured to function as an AS, each including a dedicated BGP speaker. Each simulation initialized from zero state and the BGP protocol was given 2,75 hours to reach steady state. After this initial phase, a single node was assumed to become compromised and initiate a blackhole in the network. While in this part of the simulation, the malicious agent was randomly chosen (runs were repeated 100 times), it chose an IP block for the blackhole in such way as it would (in its opinion) maximize the overall impact to the network. After the false route was advertised, we evaluated the state of the BGP tables after an additional 2,75 hour converge time.

In our analysis, we could not find any evidence for the theoretical prediction outlined in previous work. Network homogeneity does not help to reduce the impact of failures, but in fact does hurt the network in two ways: (1) As the topology structure becomes more homogeneous, the importance of the nodes to the network becomes also more homogeneous. This can be shown by the drift between the actual and expected node betweenness metric, which measures the layout of shortest paths, paths that BGP will take, in the network. In other words, as on the way from the current Internet scale-free structure, the well-connected hubs are consecutively replaced by more scale-like structures, the impact of a failure of any random node on the network grows. (2) As the hubs from the scale-free topology are exchanged, the advantage and amount of BGP routing table aggregation also disappears. A more homogeneous structure will correspondingly demand longer BGP forwarding tables across all and especially the backbone routers, which increases the demands on Internet hardware and overall forwarding delays. There exists severe concern about the recent growth of the Internet and trends such as multi-homing, load-balancing and address fragmentation, which increase the BGP routing tables sizes in the Internet by a factor of 10 [12]. Our analysis shows that moving from a pure scale-free topology to a topology in which only 25% of the network follows the homogeneous layout would increase the BGP table on average by an additional 20%. An entirely homogeneous network ($s=0$) results in a routing table close to the number of total ASes. It may be argued that while a scale-free topology limits the impact that any

random node could have on the network, the effect of a well-connected hub node will be dramatically larger. To evaluate this, we conducted a set of experiments where the compromised node was chosen according to its expected impact as measured by a high node betweenness. Our results confirm this assumption and demonstrate wide-spread outages across the network topology, but also indicate that the overall impact of the worst-case attack in a scale-free topology is only 50% worse than the worst-case in a purely homogeneous topology. The overall number of nodes that could successfully initiate such a blackholing attack is, however, very limited, and it may be assumed that these core routers would in practice have special safeguards in place to prevent such corruption.

Acknowledgements

We would like to thank the anonymous reviewers for the valuable comments. The work presented in this paper is supported by the European Commission, under Grant No. FP7-224619 (the ResumeNet project).

References

1. Callaway, D.S., Newman, M.E.J., Strogatz, S.H., Watts, D.J.: Network Robustness and Fragility: Percolation on Random Graphs. *Physical Review Letters* **85** (2000) 5468–5471
2. Cohen, R., Erez, K., ben Avraham, D., Havlin, S.: Breakdown of the Internet under intentional attack. *Physical Review Letters* **86** (2001)
3. Motter, A.E., Lai, Y.C.: Cascade-based attacks on complex networks. *Physical Review E* **66** (2002)
4. Nordström, O., Dovrolis, C.: Beware of BGP attacks. *SIGCOMM Comput. Commun. Rev.* **34**(2) (2004) 1–8
5. Rekhter, Y., Li, T., Hares, S.: RFC4271 - A Border Gateway Protocol 4 (BGP-4) <http://tools.ietf.org/html/rfc4271> (2006)
6. Zhao, X., Pei, D., Wang, L., Massey, D., Mankin, A., Wu, S.F., Zhang, L.: An Analysis of BGP Multiple Origin AS (MOAS) Conflicts. In: *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*. (2001)
7. Bono, V.J.: 7007 Explanation and Apology, <http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html> (1997)
8. : YouTube Hijacking: A RIPE NCC RIS case study, <http://www.ripe.net/news/study-youtube-hijacking.html> (2008)
9. Mao, Z.M., Govindan, R., Varghese, G., Katz, R.H.: Route Flap Damping Exacerbates Internet Routing Convergence. In: *Proceedings of the ACM SIGCOMM Conference*. (2002)
10. : (SSF-NET <http://www.ssfnet.org/>)
11. Hernandez, J.M., Kleiberg, T., Wang, H., Mieghem, P.V.: A Qualitative Comparison of Power Law Generators. In: *SPECTS*. (2007)
12. Bu, T., Gao, L., Towsley, D.: On characterizing BGP routing table growth. *Computer Networks* **45**(1) (2004) 45–54