

Analysis of BGP Prefix Origins During Google's May 2005 Outage

Tao Wan Paul C. van Oorschot

School of Computer Science
Carleton University, Ottawa, Canada
{twan, paulv}@scs.carleton.ca

Abstract

Google went down for 15 to 60 minutes around 22:10, May 07, 2005 UTC. This was explained by Google as having been caused by internal DNS misconfigurations. Another vulnerable protocol which could have caused such service outage is BGP. To pursue the latter possibility further, we explore how BGP was functioning during that period of time using the RouteViews BGP data set. Interestingly, our investigation reveals that one Autonomous System (i.e., AS174 operated by Cogent), which is apparently independent from Google, mysteriously originated routes for one of the IP prefixes assigned to Google (64.233.161.0/24) immediately prior to the service outage. As a result, 49.1% of ASes re-advertising routes for 64.233.161.0/24 switched to the incorrect path. Those poisoned ASes directly serve 1500 IP prefixes, and span a broad range of geographic locations. Since this erroneous prefix origination apparently has not occurred previously, or after this specific instance, we consider that it might have been the result of malicious activity (e.g., compromise of one or more BGP speakers) and contributed at least partially to Google's service outage.

1 Introduction

Google went down for less than an hour around 22:10, May 07, 2005 UTC [18]. One speculation is that DNS poisoning attacks caused this service outage because some traffic sent to Google was redirected to other websites for instance *sogosearch.com*, which also provide searching services [11]. Google later denied that it was under any attack and clarified that to their knowledge, their service outage was due to internal DNS misconfigurations. Regarding redirections, while we can attribute most reported cases to DNS issues (based on our communication with Google individuals – see §5.1), a few uncountered claims remain

(cf. [7] and §5.1). Nevertheless, the outage itself motivated us to examine how the Border Gateway Protocol (BGP) [15] behaved on the Internet during the period when Google was down. This is because BGP, the IETF standard and only inter-domain routing protocol used on the Internet, can be exploited to attract traffic destined to one site and redirect it to another.

We used the BGP data continually collected by the RouteViews Project [17] to analyze BGP announcements of Google's prefixes from January 1, 2005 to May 25, 2005. Interestingly, we discovered that at 14:37:56, May 07, 2005 UTC, prior to the service outage, AS174 operated by Cogent, which is apparently independent from Google, mysteriously originated routes for 64.233.161.0/24, one of the prefixes assigned to Google. This prefix contains the IP addresses associated with *www.google.com* returned from the DNS during that period of time (based on the DNS queries from a number of computers within Canada). This erroneous prefix origination did not occur prior to this specific instance, nor has it re-occurred thereafter. None of the traffic engineering approaches (e.g., multi-homing, aggregation, etc.) which we are aware of could explain this announcement. The coincidence in time with Google's service outage leads us to suggest that BGP was exploited by malicious parties to intentionally target Google.

Although we are not able to conclude definitively that the observations outlined in this note imply that the outage was the result of a deliberate BGP-related attack, we nonetheless believe it is important to present these observations publicly, to stimulate further discussion on this topic, and to highlight the other known, but not well-documented, Internet incidents that were mainly caused by misconfigurations and without any particular target.

If our speculation is indeed true, it should raise alarms that attacks exploiting routing vulnerabilities, which were forewarned about 16 years ago by Perl-

man and Bellovin [14, 2], are now reality.¹ This would strongly suggest that the Internet community should consider more seriously the design, evolution and actual deployment of security mechanisms for BGP (e.g., building on ideas presented in proposals to date including S-BGP [10], soBGP [25], psBGP [24], etc.) sooner rather than later.

This note attempts to fill a gap by documenting what may have been a real-world BGP incident specifically targeting a well-known organization. We believe it can serve the purpose of alerting the general public of the insecurity of the Internet routing infrastructure, which we hope in turn will help stimulate demand and deployment of security mechanisms for BGP.

The rest of this note is organized as follows. In section 2, we provide background information on BGP and some scenarios where multiple ASes might originate routes for a common prefix. In section 3, we describe the methodology of our analysis. Our results are presented in Section 4. We conclude in Section 5.

2 Background

We start with a brief overview of BGP, which can be safely skipped by BGP experts. We then discuss some scenarios under which the same prefix might be originated by multiple ASes at the same time (as we observed for one of Google’s prefixes); this is often referred to as Multiple Origin ASes (MOAS) [27].




Figure 1. A BGP view of the Internet

2.1 Overview of BGP

BGP is an inter-domain routing protocol based on a distance vector approach. A *BGP speaker* establishes BGP sessions over TCP with its direct neighbors, and

¹It is known that spammers commonly hijack prefixes using BGP [5]. However, we understand that they usually hijack unused address space, resulting in no harm to existing traffic flow on the Internet.

exchanges routing information with them. BGP routing information is carried in BGP UPDATE messages, which consists of Network Layer Reachability Information (NLRI) and a number of attributes (e.g., LOCAL_PREF, AS_PATH, MULTI_EXIT_DISC, etc.) associated with the NLRI. While a majority of the route attributes are not propagated beyond a neighboring AS, an AS_PATH is transitive and it consists of a sequence of ASes traversed by this route. Thus, BGP is often referred to as a *path vector* routing protocol.

We use Figure 1 to illustrate how BGP update messages get propagated across the Internet. We use $(f_x, [Y, X])$ to denote a selected portion of a BGP update message, where f_x denotes an IP prefix carried in the NLRI, and $[Y, X]$ denotes the *as_path* traversed by this route. More specifically, the route is originated for prefix f_x by AS X and has traversed through AS Y . We also say that Y readvertises the route for f_x originated by X .

Suppose IP prefix 15.0.0.0/8 (abbreviated 15/8) is assigned to a particular AS, say I . To allow other ASes to forward traffic destined to the IP addresses specified by 15/8, I originates a route $(15/8, [I])$ to its neighbors, in this case to G . After G receives $(15/8, [I])$, a series of operations are applied, including applying route importing policies, route selection procedures, and route exporting policies. If $(15/8, [I])$ passes the whole process, G will transform it to $(15/8, [G, I])$ by inserting its own AS number onto the *as_path* and then readvertises it to neighbor C .

The above process repeats at each AS in the network, which has received this route. Ideally, every AS on the network eventually receives the route originated by I . For example, J might receive $(15/8, [H, D, C, G, I])$. If every AS on the Internet originates routes for the prefixes assigned to it, the idea is that eventually every other AS may build routes for reaching the IP addresses specified by the prefixes assigned to every other AS on the Internet. In this way BGP is effective in propagating valid reachability information. However, it is also effective in propagating false routing information as we explain in §4.

2.2 Multiple Origin ASes (MOAS)

Each AS is assigned one or more IP prefixes by the organization running that AS, which either obtains the prefixes from the address authority (e.g., a RIR) or from another organization (e.g., an upstream ISP). Public data is available for looking up which organization owns a particular AS, e.g., by using one of the whois databases such as ARIN’s whois [1]. It is usually the AS to which a prefix has been legitimately assigned

which will originate a route for that prefix. In other words, there should be only one origin AS for each prefix [8]. However, some operational practices make it possible for two or more ASes to originate a route for the same prefix (referred to as MOAS). MOAS could be caused by legitimate practice (e.g., multi-homing) or by malicious attacks (e.g., prefix hijacking). Here we describe three cases of MOAS including *multi-homing*, *anycast routing*, and *prefix hijacking* (see [27] for a more detailed study).

2.2.1 Multi-homing

Many organizations connect to the Internet via two or more ISPs which may run different ASes. A multi-homing organization (X) may or may not run its own AS. In the former case, X may use a valid AS number or a private AS number. If X participates in inter-domain routing using a valid AS number, it should be the only origin AS for its prefixes. In other words, there should be no MOAS of its prefixes. If X uses a private AS number, its service providers will strip the private AS number from all routes originated by X and replace it with their own AS numbers. Thus, there will be multiple origin ASes for X 's prefixes. If X does not participate in inter-domain routing (i.e., does not run its own AS) and simply delegates its prefixes to all of its service providers, it is equivalent to the case of using a private AS number. Thus, MOAS will be observed for X 's prefixes.

2.2.2 Anycast Routing

Anycasting [13] refers to communication between a client and one of the servers within a group sharing a common IP address (anycast address). While it appears counter-intuitive for multiple servers to be configured with the same IP address, anycasting offers attractive benefits such as reduced response delay, load-balancing, and improved availability, among others. Thus, it has gained popularity among application service providers. For instance, some of the root DNS servers (e.g., F-root [9] and K-root DNS [16]) are implemented using anycast. Anycast routing refers to a practice that supports anycasting in the network layer by ensuring that a datagram sent to an anycast address is transmitted to at least one of the servers within an anycast group, likely the one "closest" to the originating network. To do so, an anycast address space will be announced by multiple routers into an interconnected network. For example, if an application service provider distributes its anycasting service across different geographic locations, each of which connects to the Internet via a different ISP, then multiple origins of

prefixes containing this anycast address space will be announced via BGP by different ASes. In other words, anycast routing can cause MOAS.

2.2.3 Prefix Hijacking

A malicious AS Y may announce a prefix assigned to another AS X without any legitimate reason, which causes MOAS. As a result, traffic originated from some part of the Internet and destined to X may be attracted to Y ; such traffic can then be manipulated in many ways. For example, traffic can be dropped; modified and then resent back to X through a tunnel; or redirected to other locations [3].

2.2.4 Aggregation

When AS Y receives an announcement of a prefix originated by X , Y may aggregate X 's prefix with other prefixes, including those assigned to Y itself. Y might appear as the origin AS of an aggregated, less specific prefix (possibly with an AS_SET or an AUTO_AGGREGATOR attribute [15]). If X also announces its prefix to another AS Z which further re-advertises the announcement, two prefix originations will occur on the Internet both of which contain the address space specified by X 's prefix. Note that prefix aggregation is not a case of MOAS, since the prefixes in question are not exactly same; we discuss it here to explain that the observed mysterious origin of one of Google's prefixes is not a result of proper aggregation, which some people might otherwise conclude is a possible explanation.

3 Methodology

We start with the discussion of the BGP data set used for our analysis, and then present our analysis results.

3.1 BGP Data Set

We used BGP data collected on a regular basis by the RouteViews Project [17] to analyze announcements of the prefixes assigned to Google. There are several BGP routers maintained by the RouteViews project, each of which establish BGP sessions with a number of ASes on the Internet. These RouteViews routers only collect BGP update messages from their neighboring ASes and do not inject any update messages back. In the absence of access to the BGP data from ASes of our choice, the BGP data collected by the RouteViews project is of central importance to our analysis. We

combine the BGP data collected by different routers to get a better view of BGP updates on the Internet. However we acknowledge that this view is still limited, and does not allow us full confidence in our conclusions on the actual reason of the Google incident, or to deduce the full impact of the incident on the Internet.

3.2 Analysis Approaches

Based on ARIN’s *whois* database [1] in June 2005, we learn that Google has AS number 15169, and is assigned two /19 and one /22 address blocks which contain in total 68 blocks of /24 prefixes. Google chooses to announce /24 prefixes instead of /19 or /22, which is a common practice for avoiding traffic destined to one AS being attracted to other ASes which might have announced that AS’s prefixes with longer prefixes. Announcing prefixes longer than 24 bears the risk of being rejected since 24 is the longest prefix acceptable to many ISPs. Based on the BGP data we used, AS15169 has 25 neighbors including AS174.

Our hypothesis is that if someone tried to attract traffic destined to Google by prefix hijacking, we should see MOAS regarding some of Google’s prefixes. Thus we first looked at one BGP routing information base (RIB) collected near the time when Google went down. We discovered one AS (i.e., AS174) in fact originated 64.233.161.0/24, one of the prefixes assigned to Google, before Google’s outage. We then analyzed the RIBs collected over a number of days to determine the duration of this mysterious announcement, or what we call the *MOAS period*. We then analyzed one BGP RIB per day from January 1, 2005 to the start of the MOAS period, namely in the *pre-MOAS* period, and one BGP RIB per day from the end of the MOAS period to May 25, 2005, namely in the *post-MOAS* period. We then compared how AS174 originated routes for the prefixes assigned to Google over these periods.

4 Results of Analysis

Here we first report our direct observations from the RouteViews data regarding the advertisements of 64.233.161.0/24, and then give our interpretation of the results.

4.1 Direct Observation

We observed that AS174 started to originate Google’s prefix from 14:37:56, May 07, 2005 UTC and stopped after 10:52:00, May 09, 2005 UTC; we call this the *MOAS period*. We next report our observations

respectively for the three periods (pre-MOAS, during-MOAS, and post-MOAS period). We repeat that all of our observations are from the vantage point of the RouteViews, unless stated explicitly otherwise. However, this limited view has no impact on our observed fact that AS174 originated one of Google’s prefixes.

4.1.1 Pre-MOAS Period

Prior to the MOAS period, we observed that AS15169 originated 64.233.161.0/24 to 25 direct neighbors including AS174, which further re-advertise the route to 31 remote ASes (i.e., not directly connected to AS15169). In total, we observed that 56 ASes re-advertised routes for 64.233.161.0/24, and that no AS other than AS174 itself re-advertised routes with an AS_PATH involving AS174 (hereafter “via AS174”). In other words, we did not observe any AS re-advertising routes for 64.233.161.0/24 via AS174 (see Figure 2). Thus, it is very likely that in the pre-MOAS period, traffic destined to 64.233.161.0/24 passed through AS174 only if the traffic originated from AS174 or from its customers.

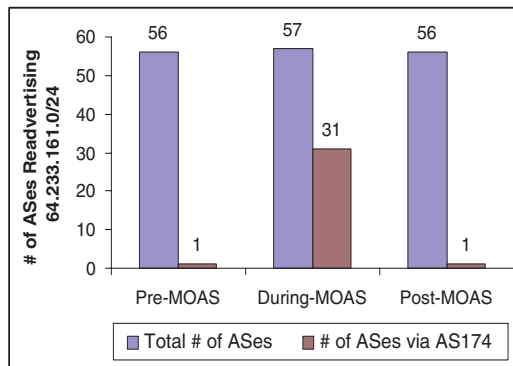


Figure 2. Total number of ASes re-advertising routes for 64.233.161.0/24, and total number of ASes via AS174

4.1.2 During MOAS Period

During the MOAS period, AS174 originated routes for 64.233.161.0/24 instead of re-advertising the one originated by AS15169, thus poisoning many ASes’ routing tables. We observed in total 57 ASes re-advertising the routes for 64.233.161.0/24, among which 31 ASes preferred (i.e. selected) the routes originated by AS174. Among these 31 ASes, 28 of them switched from the routes originated by AS15169 to those originated by

AS174, including 8 of AS15169’s direct neighbors. We refer to these as *poisoned* ASes. Some of the poisoned ASes are large ISPs, such as AS701 (UUNET), AS2497 (IJJ), AS3561 (C&W), and AS7018 (AT&T). Geographically, they span almost every continent. In terms of percentage, 49.1% (28 out of 57) of re-advertising ASes were poisoned, including 32% of AS15169’s direct neighbors (see Figure 2).

We examined the prefixes assigned to the poisoned ASes for perspective on the amount of address space from which traffic originated toward Google might have been attracted to AS174. Based on the data we used, in total 2003 prefixes were assigned to the 28 poisoned ASes. Figure 4 presents those prefixes arranged by prefix length. This demonstrates that not only prefixes containing relatively small address ranges (e.g., /24) are affected, but also some prefixes containing larger address space (e.g., with a length shorter than 16). This is not a surprise since some of the poisoned ASes are large ISPs which hold a large amount of address space.

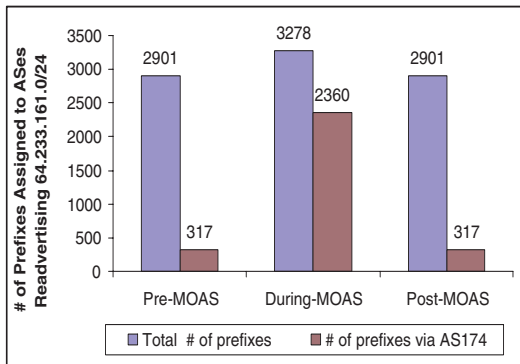


Figure 3. Total number of prefixes assigned to the ASes re-advertising routes for 64.233.161.0/24, and the number of prefixes assigned to poisoned ASes

4.1.3 Post-MOAS Period

After the MOAS period, we observed in total 56 ASes re-advertising the route for 64.233.161.0/24 originated by AS15169. No one among the 56 observed ASes except AS174 re-advertised 64.233.161.0/24 with an AS_PATH involving AS174. This is the same situation as in the pre-MOAS period.

Regarding other prefixes assigned to Google, we did not observe any multiple origins for any of these, neither by AS174 nor any other ASes during the three periods.

4.2 Our Interpretation

First, to our knowledge, none of the legitimate reasons as discussed in §2.2 can explain why AS174 would originate the IP prefix assigned to Google. We acknowledge that there may be many business-related issues beyond our knowledge, which may affect BGP operational practice, as suggested by a recently reported disagreement [26] between previous BGP peers, Cogent and Level3. To this end, we next consider the possibility that this incident was caused by misconfiguration or malicious attack.

4.2.1 Misconfiguration

We consider two types of misconfiguration which might result in the MOAS regarding 64.233.161.0/24. Firstly, many ASes use centralized databases, which contain IP prefixes assigned to an AS, to automatically generate configuration files for BGP speakers within an AS. If a prefix f_x assigned to AS X erroneously enters into the central database from which AS Y draws its BGP speaker configurations, AS Y might erroneously originate routes for f_x . So it is possible in theory but unlikely in practice that a single prefix 64.233.161.0/24 got into AS174’s configuration database by mistake and AS174 updated some of its BGP speakers using the misconfigured database before the MOAS period. Secondly, it is also possible that one or more BGP speakers in AS174 were misconfigured such that they stripped the origin AS from a route when re-advertising that route. However, this second situation appears very unlikely since we did not observe the same misbehavior happening on any other prefixes announced by AS15169 to AS174.

4.2.2 Malicious Attack

It is also possible that one or more of the BGP speakers in AS174 were compromised and used to influence traffic sent to Google. While some traffic destined to 64.233.161.0/24 indeed was forwarded through AS174 to AS15169, and an attacker with control of a BGP speaker in AS174 could get access to that traffic without hijacking Google’s prefix, the amount of such accessible traffic is limited and a large portion was forwarded to AS15169 by its neighbors other than AS174. Thus, hijacking the prefix allows an attacker to gain access to more traffic destined to the hijacked address space. An attacker would have considerable freedom in manipulating the attracted traffic, depending on how much control he had over the compromised routers. A simple attack is to redirect traffic to a black hole by installing unreachable static routes in the routing




Figure 4. # of Prefixes with Specified Length Assigned to the Poisoned ASes

table of the compromised router in AS174. An advanced attack is to redirect attracted traffic to a location (e.g., a compromised PC) where their destination IP addresses are replaced by new IP addresses (e.g., the IP addresses of other websites). The modified traffic is then re-injected into the Internet [4, 3]. If an attacker chose to not manipulate the attracted traffic, the traffic might still be able to reach its intended ultimate destination, i.e., Google, since AS174 has direct connectivity to AS15169.

5 Communication with Google and Cogent

We made attempts to obtain inputs from both parties involved in this incident, i.e., Google and Cogent. Here we summarize our communication with them.

5.1 Communications with Google

Our communication with Google [22] was useful on several fronts. First, we acquired better understanding of Google’s internal DNS failures which led to the outage. The failure was caused by an unreadable configuration file consisting of almost all DNS A records that was mistakenly pushed to all Google DNS servers. As a result, all DNS queries sent to Google’s DNS servers were returned with answers of no A records, which in turn caused the service outage. We also learned that Google does not use anycast routing, which otherwise might have provided a logical explanation for the MOAS which occurred.

Second, we understood that most of the reported traffic redirections were mainly caused by browsers trying to append a Top Level Domain or TLD (e.g., .com) when a supplied domain name could not be resolved. In this case, many queries intended to *google.com* ended

up at sites such as *google.com.com*, which happens to be hosted by *sogosearch.com*. In addition, some Internet Service Providers (ISPs) return pointers to some special sites whenever a domain name search fails, which might have also caused some redirections.

Third, we learned that our observed MOAS is not considered legitimate by Google for their prefixes, and Google did experience problems with 64.233.161.0/24 during the period of MOAS at the Point of Presence (PoP) through which Google peers with Cogent. However, statistical traffic differences were not apparent given the large volume of traffic received by Google. We also came to agreement that a few uncountered claims [7], could indeed have been caused by redirections involving BGP. For example, some traffic sent to “*www.google.com*” was redirected to “*search.msn.com*” (cf. [7]). Such redirection appears unlikely to have been caused by attempts to append a TLD.

Fourth, our draft report served the purpose of alerting Google personnel to BGP security issues. After reading our draft, we were told that Google’s network operation group was “sufficiently disturbed” by the fact that BGP can be used for prefix hijacking to consider setting up infrastructure for monitoring apparent hijacking of Google’s IP prefixes.

5.2 Communications with Cogent

We have made several attempts to discuss this report with individuals from Cogent. We first contacted Cogent Network Operation Center (NOC) at “*noc@cogentco.com*” [20]. We were asked for the AS_PATHs involved in the incident, and our relationship with Google. After providing the requested information, we did not hear back further.

Our second attempt involved sending a request [23] to the NANOG mailing list, asking for a technical contact at Cogent to discuss BGP issues. Our email to

the NANOG mailing list resulted in email exchange [19] with an employee from a Cogent help-desk who advised us that she/he was not able to discuss this report with us due to privacy agreements. A separate email contact [21] through Cogent NOC proved to be equally unhelpful. While this does not provide any evidence supporting our conjectures, neither does it contradict any.

While we continue to welcome input from Cogent, in the absence of further ideas on how to confirm or deny our conjectures, we feel it is in the best interest of the community to make our report available for others (who might have access to more information than us) to draw their own conclusions.

6 Concluding Remarks

While some MOAS is valid, we (the authors) are not able to find any technical explanation for the observed event other than one or more BGP speakers within AS174 having misbehaved. On the extreme, any such misbehaving BGP speakers might have been controlled by an attacker which then redirected Google's traffic to other sites of the attacker's choice.

This Google outage incident differs from other BGP incidents [12, 6] in that it is subtle and might have involved malicious activity specifically targeting an organization, while others are known to have been caused by misconfiguration and without any specific target. This incident, among others, again highlights that BGP is extremely vulnerable and need be secured to protect the Internet, which is now clearly recognized as a critical infrastructure and is on the path to replace many of the traditional communication infrastructures (e.g., telephony networks).

We now comment briefly on the difficulty in making real-world progress toward securing BGP, which requires collaboration among many parties, e.g., router vendors and ISPs. While many stake holders are aware of the problem, none has taken initiative to push it forward. One operational obstacle is that extra costs will be incurred as a result of developing and deploying BGP security solutions. With the current downturn in the telecommunications industry, cost reduction has become a primary objective of many router vendors and ISPs. Thus, it appears unrealistic to expect ISPs to start to spend on deploying BGP security solutions which do not provide to them an immediate return on investment unless there is a strong demand from their customers. Router vendors likewise appear unmotivated to develop BGP security solutions, due to the lack of interest from ISPs.

We suggest that governments can play an impor-

tant role to facilitate the development and deployment of more secure versions of BGP. While the Internet is mainly built and operated by ISPs, it is now of general public interest since many people and the majority of businesses are reliant on the Internet for their daily activities. Thus, we believe that it is a natural role for top-level governments – and one few if any other parties may take on responsibility for – to ensure that the Internet in general, and BGP in particular, is secured, especially from a robustness and survivability perspective. As a tangible example, governments could provide additional funding to stimulate research and development of BGP security solutions; might encourage ISPs to deploy BGP security solutions (e.g., by subsidies or tax credits or other incentives); or maybe even require that the Internet routing infrastructure used within the government itself must employ a more secure version of BGP. Such requirement may be particularly effective, because of the very significant spending power of very large governments, and their resulting economic leverage over vendors of Internet infrastructure services.

Acknowledgements

We thank Steve Bellovin, Stephen Chou, Evangelos Kranakis, Marcus Leech, David Presotto, and anonymous reviewers for their comments on a draft of this paper, which helped us improve its present form. Any remaining errors, including possible mis-interpretation of the data we have observed, remain the responsibility of the authors. The first author is supported in part by MITACS (Mathematics of Information Technology and Complex Systems) and NCIT (National Capital Institute of Telecommunications). The second author is Canada Research Chair in Network and Software Security, and is supported in part by the Canada Research Chairs Program, an NSERC Discovery Grant, and MITACS.

References

- [1] ARIN. American Registry of Internet Numbers, 2005. <http://www.arin.net>.
- [2] S. Bellovin. Security Problems in the TCP/IP Protocol Suite. April 1989.
- [3] S. Bellovin. Routing Security. British Columbia Institute of Technology, June 2003.
- [4] S. Bellovin. Where the Wild Things Are: BGP Threats. NANOG 28, June 2003.

- [5] S.M. Bellovin. Spamming, Phishing, Authentication, and Privacy. *Communications of the ACM (Inside Risks)*, 47(12), December 2004.
- [6] L. Blunk. New BGP Analysis Tools and a Look at the AS9121 Incident, March 2005. <http://www.merit.edu/nrd/papers-presentations/MTP-2005-01.pdf>.
- [7] P. Ferguson. Google DNS Problems, May 2005. <http://www.merit.edu/mail.archives/nanog/2005-05/msg00238.html>.
- [8] J. Hawkinson and T. Bates. Guidelines for Creation, Selection, and Registration of an Autonomous System (AS). IETF RFC1930, March 1996.
- [9] ISC. F-ROOT. <http://www.isc.org/ops/f-root/>.
- [10] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4), April 2000.
- [11] O. Malik. Om Malik's Broadband Blog, May 2005. <http://www.gigaom.com/2005/05/07/google-hacked/>.
- [12] S. A. Misel. Wow, AS7007! NANOG mail archives, <http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html>.
- [13] C. Partridge, T. Mendez, and W. Milliken. Host Anycasting Service. RFC 1546, November 1993.
- [14] R. Perlman. Network Layer Protocols with Byzantine Robustness. Technical Report MIT/LCS/TR-429, October 1988.
- [15] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP 4). IETF RFC1771, March 1995.
- [16] RIPE. K-ROOT. <http://k.root-servers.org/>.
- [17] RouteViews. Route Views Project, 2005. <http://www.routeviews.org>.
- [18] I. Thomson. Possible DNS Hack Knocks out Google, May 2005. <http://www.vnunet.com/news/1162902>.
- [19] T. Wan. Personal Communication with Cogent Help Desk, September 2005.
- [20] T. Wan. Personal Communication with Cogent Network Operation Center (NOC), July 2005.
- [21] T. Wan. Personal Communication with Cogent Network Operation Center (NOC), September 2005.
- [22] T. Wan. Personal Communication with David Presotto of Google, June 2005.
- [23] T. Wan. Technical Contact at Cogent, September 2005. <http://www.atm.tut.fi/list-archive/nanog-2005/msg07310.html>.
- [24] T. Wan, E. Kranakis, and P.C. van Oorschot. Pretty Secure BGP (psBGP). In *Internet Society Proceedings of the Symposium on Network and Distributed Systems Security (NDSS'05)*, 2005.
- [25] Russ White. Securing BGP Through Secure Origin BGP. *The Internet Protocol Journal*, 6(3):15–22, September 2003.
- [26] C. Wilson. Peers or Not? Cogent, Level 3 Disagree, October 2005. http://telephonyonline.com/home/news/cogent_level_3_100505/.
- [27] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang. An Analysis of BGP Multiple Origin AS (MOAS) Conflicts. Proceedings of ACM SIGCOMM Internet Measurement Workshop, 2001.