Analysis of BGP Update Surge during Slammer Worm Attack ⋆

Mohit Lad[1], Xiaoliang Zhao[2], Beichuan Zhang[2], Dan Massey[2], and Lixia Zhang[1]

[1] University of California, Los Angeles, CA 90025, USA
[2] USC Information Science Institute, Arlington, VA 22203, USA

Abstract. Although the Internet routing infrastructure was not a direct target of the January 2003 Slammer worm attack, the worm attack coincided in time with a large, globally observed increase in the number of BGP routing update messages. Our analysis shows that the current global routing protocol BGP allows local connectivity dynamics to propagate globally. As a result, any small number of edge networks can potentially cause wide-scale routing overload. For example, two small edges ASes, which announced less than 0.25% of BGP routing table entries, contributed over 6% of total update messages observed at monitoring points during the worm attack. Although BGP route flap damping has been proposed to eliminate such undesirable global consequences of edge instability, our analysis shows that damping has not been fully deployed even within the Internet core. Our simulation further reveals that partial deployment of BGP damping not only has limited effect, but may also worsen the routing performance under certain topological conditions. The results show that it remains a research challenge to design a routing protocol that can prevent local dynamics from triggering global messages in order to scale well in a large, dynamic environment.

1 Introduction

The SQL Slammer worm [1] was released on Jan 25th, 2003 and exploited a known bug in MS SQL servers. Infected machines sent heavy traffic loads towards seemingly random destinations. If a destination shared the same MS SQL vulnerability, it would become infected and in turn attempt to infect another set of random destinations. Slammer infected at least 75,000 hosts in just over 30 minutes and is reported to be the fastest spreading worm to date [2]. Although the SQL Slammer worm attack was *not* directly targeted at the Internet routing infrastructure, the Internet Health Report [3] reported that a number of critical AS-AS peering links were operating above critical load thresholds during the Slammer attack period. Further, [2, 4, 5] noted that the Slammer worm attack coincided in time with a large, globally observed increase in the number of BGP[6] routing update messages. In fact, such coincidences between Internet worm attacks and the surges in BGP routing update messages have been also observed during

⋆ This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No DABT63-00-C-1027 and by National Science Fundation(NSF) under Contract No ANI-0221453. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the DARPA or NSF.

previous worm attacks, such as the Code-Red [7] attack in July 2001 and the NIMDA [8] attack in September 2001.

In this paper, we analyze the BGP log data collected from various monitoring points to understand the causes of the high surge in BGP update messages during the SQL Slammer attack. Our analysis shows that the current BGP routing protocol allows *local* connectivity dynamics to propagate *globally*. As a result, any small number of edge networks, such as those networks whose connectivity to the Internet was severely impacted by Slammer, can potentially cause global routing overload. Figure 1 shows the number of BGP path announce messages as observed from routers in three different ASes. The results, typical of that seen by monitoring points located in other ASes, clearly show a surge in BGP activity that coincides with the worm attack. As we will show later in the paper, two small edges ASes, which announced less than 0.25% of BGP routing table entries, contributed over 6% of total update messages observed, during the worm attack.

BGP route flap damping[9] was introduced specifically to prevent edge instability from flooding update messages globally. Route damping is applied on a $\langle peer, prefix \rangle$ basis and each update received from peer N for prefix p increases a penalty value associated with $\langle N, p \rangle$. Once the penalty exceeds a threshold value, N's updates regarding p are suppressed and the router behaves as if N has no route to p. The penalty decreases (decays) exponentially using a configured half-life and routes from N are again accepted after the penalty falls below a re-use limit. The RFC suggests that all Internet core routers deploy damping in an effort to "provide a mechanism capable of reducing router processing load caused by instability and in doing so, prevent sustained route oscillations". However, we will show that BGP route flap damping has not been fully deployed within the Internet core. Furthermore, even if a full deployment of route flap damping could reduce the number of global updates, it also results in a much longer routing convergence time as shown by both [10] and our simulations. Our simulation further reveals that partial deployment of BGP damping not only has limited effect but may also *worsen* the routing performance under certain topological conditions.

The paper is organized as follows. Section 2 presents our study of BGP behavior during the Slammer worm and demonstrates how local changes can propagate to create global events. Section 3 examines the impact of BGP route flap damping on the observed data and shows evidence that route flap damping could have reduced some dynamics if it had been deployed. Section 4 uses simulation to explore the impact of route flap damping deployment on both update counts and convergence time. Section 5 reviews the related work and Section 6 concludes the paper.

2 Edge AS Instabilities With Global Effects

[2, 4, 5] noted that the Slammer worm attack coincided in time with a large increase in BGP routing update messages. To understand the worm's effect on BGP, we delve deeper into the update behavior during the SQL Slammer worm attack, with the objective of investigating the origins of the update bursts and their time distribution. Our results show that a small number of edge AS prefixes contribute greatly to the global routing update volume.

2.1 Methodology

Our study uses BGP update data collected by Oregon RouteViews[11]. Monitoring points at RouteViews peer with routers in a wide variety of Autonomous Systems. Each AS router treats the monitoring point as a BGP peer router and sends routing updates to the monitoring point, where they are logged and made available to researchers. We examined BGP update logs from January 22, 2003 to January 29, 2003, a period of 8 days around the Slammer worm attack. We apply the techniques from [12] to remove monitoring artifacts from the logs.

The resulting data shows a dramatic increase in the number of BGP route announcements and matches the observations reported in [2, 4, 5]. Figure 1(a) shows the number of route announcements sent by monitored routers in three difference Autonomous Systems. Early on January 25, the number of update spikes dramatically and tails off by January 26. Similar large surges in BGP updates are seen at all monitored AS. Figure 1(b) further shows the presence of a spike in the number of withdraw messages observed during the same worm period. This suggests that there was also loss of connectivity to some portions of the Internet.

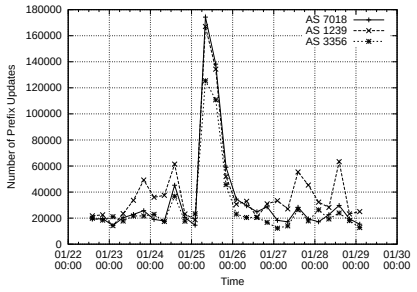

(a) Path Announcements from January 22, 2003 to January 29, 2003

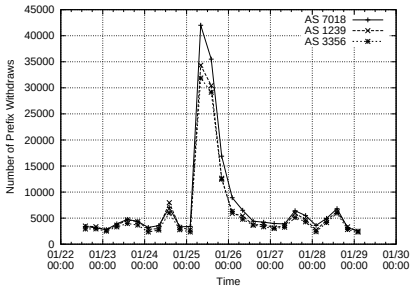

(b) Withdrawals from three peers from January 22, 2003 to January 29, 2003

Fig. 1. Updates Messages During Slammer Attack

2.2 Identifying Edge Instability

[2] reported that the worm caused high traffic congestion on the edge ASs. To better understand the edge behavior, we ranked the Internet ASs based on the average number of update messages per prefix originated by the AS, as observed on January 25 . Figure 2(a) shows the top five ranked ASs. Figure 2(b) shows the corresponding total updates for prefixes originating from these AS, as observed from all the monitoring points. For comparison, Figures 2(c) and 2(d) show the corresponding graphs on the day prior to the worm attack; note the maximum value on y-axis decreases by two orders of magnitude.

Figure 2(a) shows that prefixes from AS 18296 (belonging to a unversity in South Korea), accounted for the highest number of updates per prefix, averaging over 4500

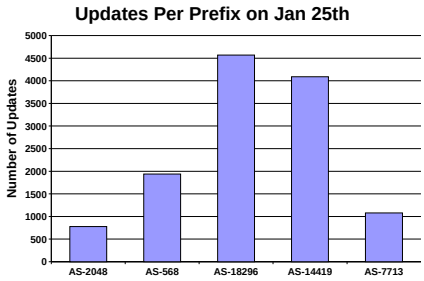

(a) Average Updates per Prefix on Jan 25, 2003

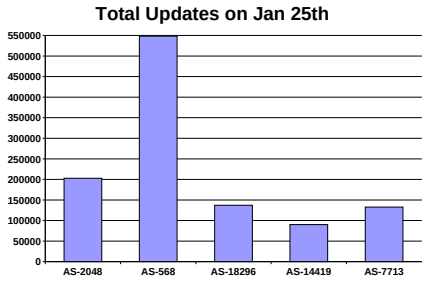

(b) Total Updates on Jan 25, 2003

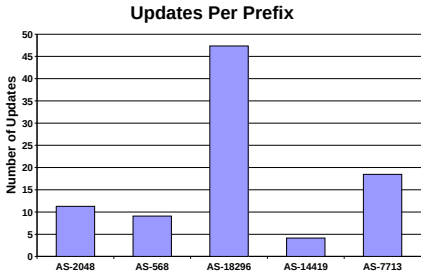

(c) Average Updates per Prefix on Jan 24, 2003

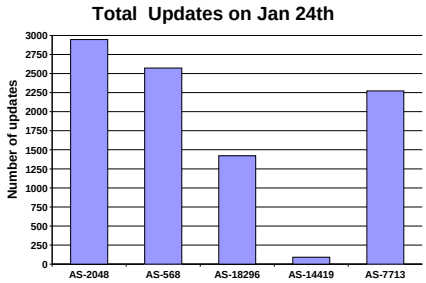

(d) Total Updates on Jan 24, 2003

Fig. 2. Update counts during Slammer Worm attack

updates per prefix on the day of Slammer attack. In contrast, this same AS, averaged only 47 updates per prefix on the day prior to the Slammer attack. This massive change is not suprising, following reports in [13], suggesting that South Korea's connectivity was among the worst affected by the SQL Slammer worm. AS 18296 advertises only about 30 prefixes out of the roughly 120,000 prefixes in a typical global BGP routing table. Although AS 18296's 30 prefixes constitute less than 0.02% of the total prefix space, this AS generated about 1.7% of the total BGP updates observed on January 25. AS568, owned by the US Department of Defense, advertises a total of 238 prefixes and stands out in terms of total number of updates. Together AS 18296 and 5568 announce less than 0.25% of Internet prefixes, but contributed over 6% of total updates seen during the worm attack, as shown in Figure 3(a).

2.3 Analysis of AS 18296 and AS 568 Edge Instability

Figure 3(b) classifies the BGP update messages associated with the 30 prefixes in AS18296 as viewed from a particular ISP, AT&T (AS7018). The updates in Figure 3(b) are classified into five categories: *DPATH* updates indicate a change in the AS path;

New Announcement updates announce a path to a previously unreachable prefix; *Withdrawal* updates remove the path to prefix and make the prefix unreachable; *Duplicates* convey no new information what so ever and are identical to the previous update for this prefix; finally, *SPATH* (Same AS Path) updates indicate no change the AS path, but do indicate a change in some other BGP attribute. On the worm attack day, DPATH messages (Different AS Path messages) were the dominant type of BGP update associated with the AS 18296 prefixes; the number of withdraw messages is also significant. Both DPATH and withdraw BGP updates convey real changes in the AS paths used to reach these prefixes. Similar results are obtained when the AS 18296 updates are analyzed from other ISPs and we note that these prefixes are seen as unstable from a wide variety of locations, an indication that the cause of the instability is close to the origin AS.

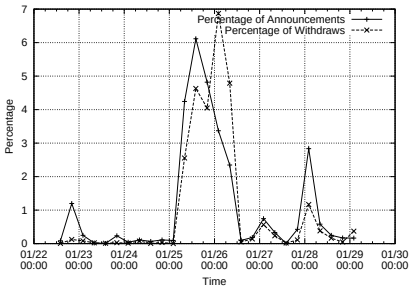

(a) Percentage of updates on prefixes belonging to AS18926 and AS568 from Jan 22nd to Jan28th

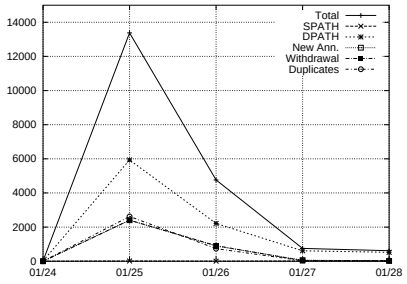

(b) Classification of Update Messages for AS 18296

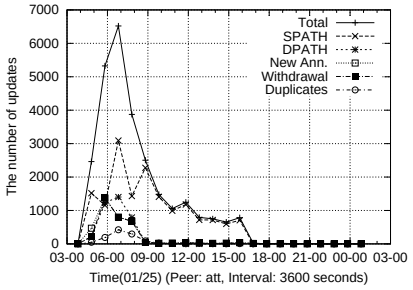

(c) Classification of Update Messages for AS 568

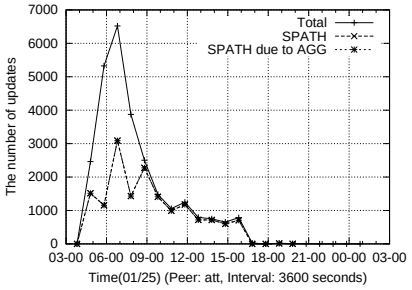

(d) Aggregator Attribute for AS 568 Updates

Fig. 3. Statisticcs for AS 568 and AS18296

AS 18296 primarily uses AS 9318 as a next hop to reach the Internet, although the data shows it is also multi-homed through AS 4766. Both prior to and after the Slammer attack, AS 18296 announced only a small number of BGP updates each day. A routing snapshot from Jan 22, 2003 showed that nearly all of the prefixes originating from AS

18296 had a next hop of AS 9318. We also observed that a total of 220 prefixes relied on AS 9318, thus AS 18296 originates roughly 14% (30/220) of the prefixes that rely on AS 9318. During the worm attack day, however, the AS 18296 prefixes accounted for 82% of BGP updates involving AS 9318. Every monitoring point observed high numbers of updates for the 30 prefixes originated by AS 18296, indicating a problem near the origin. Furthermore, the AS next to the origin, AS 9318, exhibited few changes in connectivity to destinations except for the prefixes originating from AS 18296. The above evidence strongly suggests that the problem is local to AS 18296, or the peering between AS18296 and AS9318. We can see from this example, that the worm attack affected edge ASs, causing fluctuations in connectivity, and thus resulting in a burst of BGP updates.

Figure 3(c) shows the type of BGP update messages generated by AS 568's 281 prefixes. Our previous work in [14] demonstrated that congestion at AS 568 coincides with a local AGGREGATOR attribute changes and the BGP design propagates this local change globally. Figure 3(d) shows that nearly all of the SPATH updates reported a change in the AGGREGATOR attribute.

3 Worm Attack and BGP Damping

Section 2 showed that the worm caused a high number of updates and identified particular edge AS instability that was partly responsible for the update surge. This section investigates how BGP features designed to limit instability, fared during the worm attack. BGP includes a Minimum Route Advertisement Interval (MRAI) timer that places a lower limit on the time between two successive update messages The MRAI timer is intended to suppress routing instability during this hold time, and help BGP convergence by limiting AS path exploration [15]. Fig. 4 shows the per-prefix cumulative inter-arrival time between BGP update messages, as observed from AS 7018. We can see that about 50% of the update messages have an inter-arrival time per prefix of close to MRAI timer's default value, 30 seconds. While the update surge could have been worse without MRAI, the presence of MRAI timer alone was clearly not sufficient to prevent a sudden update burst. *BGP route flap damping* is also intended limit the impact of edge instability.

3.1 Route Flap Damping

While the MRAI timer works at the time scale of tens of seconds, *route flap damping* [9] was proposed to deal with route instability at a larger time scale. The objective of damping is to suppress the updates caused by an unstable link from propagating. Unlike MRAI timer operation that is implemented on the sender side, damping is implemented on the receiving end of any peering session. A BGP router maintains a penalty value for each prefix and peer combination it receives. Whenever a peer advertises a change in route to the prefix, its penalty is increased. When a route's penalty exceeds a *suppression threshold*, a BGP router will stop using this route, thus preventing future changes from propagating. The penalty decays exponentially over time and a suppressed route

can be reused only after the penalty over time, drops below a *reuse threshold*. Table 5 shows the default damping settings for Cisco routers and Juniper routers

[9] recommends that at least the Internet core routers enable damping to achieve reduced update messages. Damping has been said [16] to be widely deployed and is considered a major factor in keeping the Internet stable at BGP's early days, when faulty implementations caused lots of excessive updates.

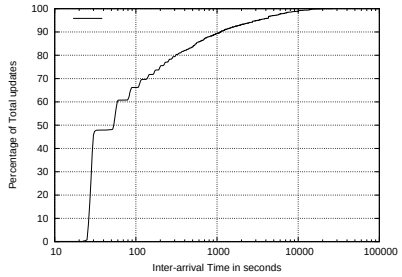

Fig. 4. Cumulative Distribution of Inter-arrival time for AS 18296 for Jan 25th 2003 from AS 7018

Damping Parameter	Cisco	Juniper
Withdraw penalty	1000	1000
Re-announcement penalty	0	1000
Attributes change penalty	500	500
Suppression threshold	2000	3000
Half time (min)	15	15
Reuse threshold	750	750
Maximum suppress time (min)	60	60

Fig. 5. Default parameters for route flap damping

3.2 Case Study on Damping

We used the monitoring point data to infer BGP activity on real operational links. Due to the limitations of BGP monitoring and data collection, some assumptions are needed to make the analysis feasible. As in most BGP research work, we model the Internet inter-AS topology as a graph, in which each AS is represented by a single node, and every AS link appearing in any update message is represented by a single link between the two AS nodes. In our analysis, we choose sequences of updates that exhibit simple route flap behavior, such as alternating between two paths, or up and down for a single path, so that we can safely infer remote BGP activity without making broad assumptions on routing policy. One such sequence of frequent updates was received from AS 7018 for the path to AS 18296, and we analyzed it to see if route flap damping should have been triggered. Fig. 6(a) shows the topology containing the relevant ASes and links. The link we are interested is the one connecting AS 7018 and AS 701. Table 1 shows the first few updates in this sequence. The penalty is calculated using Cisco default parameters.

We assume that the first withdraw creates a consistent routing state (i.e., no one in Fig. 6(a) has a path to AS 18296) and start our analysis from this point, with the penalty value set to 0. The second update announces a path via AS 701, which can only happen if AS 701 has sent such an announcement to AS 7018. Since it is a re-announcement following a withdraw, the damping penalty is zero. The third update shows that AS 7018 changes its next hop to AS 3549, and the fourth update withdraws the path again. The explanation of these two updates depends on AS 7018's routing policy. If AS 7018 prefers AS 701 over AS 3549, it could be that AS 701 withdraws its path first, causing AS 7018 to switch to AS 3549 and until AS 3549 withdraws its path as well. If AS 3549

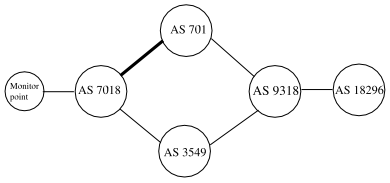

(a) Topology in Case 1

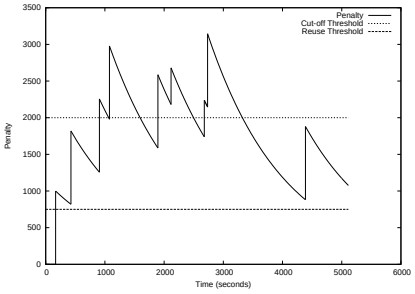

(b) Damping Penalty on link between AS 7018 and AS 701

Fig. 6. Case study of peering between AS7018 and AS701

is more preferred, the third update could be caused by an announcement from AS 3549, followed withdraws from both AS 701 and AS 3549 that cause the fourth update. In our analysis, instead of estimating an AS's routing policy, we simply calculate penalty for all possible cases and choose the one that gives the most conservative value.

No.	Time (second)	Update observed	Possible update on interested link	Induced Penalty	Total Penalty
1	0	withdraw	withdraw	-	0
2	30	(7018 701 9318 18296)	re-announcement	0	0.0
3	141	(7018 3549 9318 18296)	-	0	0.0
4	167	withdraw	withdraw	1000	1000.0
5	281	(7018 701 9318 18296)	re-announcement	0	915.9

Table 1. Sample update sequence and analysis

Fig. 6(b) shows the damping penalty value on link between AS 7018 and AS 701 over a time period of about 80 minutes. The penalty was well above the suppression threshold, reaching more than 3000 while the suppression threshold is only 2000. However, we still see updates from AS 7018, inspite of calculations showing that the threshold was crossed. Had damping been turned on, update from AS 7018 should not have included AS 701 until the penalty dropped below the reuse threshold and this may have prevented the unstable path information from spreading to other parts of the Internet. Therefore, this case study suggests that, even in the Internet core, it is possible that AS 7018 did not implement route flap damping at its link to AS 701. Overall, our analysis suggests that the ideal core deployment has not been achieved in practice.

4 Effectiveness of Route Damping

BGP route flap damping could have an impact on the number of updates, but the extent of damping is unclear based on data observed. In this section, we use simulations to

evaluate the potential effectiveness of route flap damping in response to instability at an edge AS, such as that seen during the SQL worm attack.

4.1 Simulation Settings

We used the SSFNET BGP simulator [17] to simulate BGP behavior with different topologies. In our simulations, we set link delay to 2 milliseconds and the processing delay of each routing message to a random value between 0.1 and 0.5 second. The MRAI timer (discussed in the previous section) is set to the default value of 30 seconds with random jitter, the value used widely in real network operations. We used 110-node and 208-node topologies that are derived from real Internet routing tables [18].

For each simulation run, we randomly chose one node from the topology and attached an extra origin AS node to it. Throughout the rest of the paper we refer this extra AS node as the *flapping source*. To simulate edge instability, the flapping source withdraws its AS path every 100 seconds, and re-announces it 50 seconds later. We also simulated different rates, but the results were found to be similar. A pair of the withdraw and re-announcement is called a *pulse*. For each unstable source, we ran the simulation using one to 25 pulses. The route flapping starts after the network has been stable for a 1000 second period. Each simulation run ends when there are no more updates to send, *i.e.*, when BGP converges. We count the number of update messages during this time period in order to evaluate the effective damping has on reducing updates.

4.2 Results

Figure 7 shows the results for three different route damping deployment scenarios:

1. No damping deployed at any node
2. Full deployment of damping at every node
3. Damping deployment only at “core” nodes

In our simulation, a ”core” node is distinguished by its node degree and nodes with degree larger than or equal to 13 as core nodes. 8 (7.3%) nodes were selected as core nodes. The X-axis represents the number of pulses generated by the instable edge AS (flapping source) and for each number of pulses, we run a set of simulations by randomizing the location of flapping source. The Y-axis represents the total number of updates in the network during the simulation. The results presented are the average values with 95% confidence interval.

The results show that full deployment of route flap damping dramatically reduces the total number of updates and the reduction in updates occurs regardless of the location of flapping source. In Figures 7(a) and 7(b), the number of updates sent with no damping deployed increases linearly as the number of flaps increases. However, with deployment of damping at every node, the number of updates initially grows and then remains nearly constant. The neighbors immediately adjacent to the flapping source will damp the prefix from the unstable edge AS and thus effectively ignore the updates caused by additional flaps, preventing these updates from propagating throughout the network. This can help reduce the routers’ load when excessive updates were generated

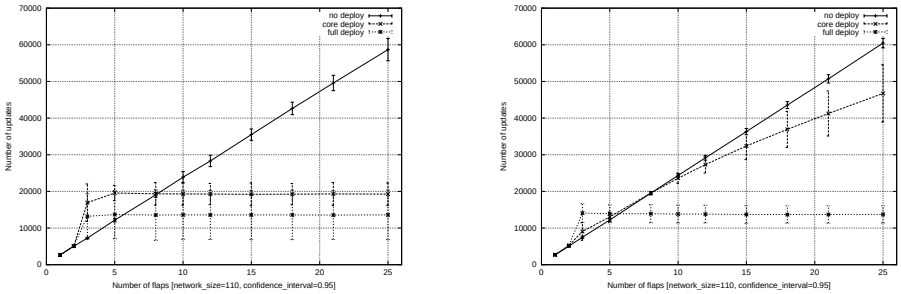

(a) Flapping source attached to core

(b) Flapping source attached to edge

Fig. 7. Number of Updates

by highly intermittent connectivities, or stressful network events, such as worm attacks. In addition, our simulations (not shown in this paper) also confirm the findings in [10], that route flap damping can result in a much longer route convergence time.

However, even deployment of full damping raises issues that are not yet well understood. Note that for a small number of initial flaps (less than 8 for flapping source attached to the core and less than 5 for flapping source attached to an edge), the number of updates actually increases due to route damping. This behavior reflects the complexity of adding features to a complex distributed system. In this case, adding a route suppression feature actually *increases* the total update count for the system in the special case of a small number of origin AS flaps.

When damping is deployed only at core nodes as the BGP standard recommends, our simulation results show that damping has different effects depending on whether the flapping node is connected to a core node or an edge (non-core) node. As shown in Figure 7(a), when the flapping source was attached to core nodes, the damping effectively reduced the number of updates. In this case, the node directly next to the flapping source would damp future changes after a sufficient number of initial flaps. However, when flapping source was attached to a non-core node, the damping was not as effective in reducing the number of updates, as shown in Figure 7(a). In both the core and non-core cases, the deployment of damping will still incur a cost of delayed convergence.

In conclusion, while damping appears to be a useful technique to suppress updates due to unstable links, there is a lack of complete understanding of the effects of the deployment of damping. On one hand is the issue of increased convergence time, while on the other hand as showed in this paper, further work needs to be done on understanding the tradeoffs of partial deployment.

5 Background and Related Work

The Internet consists of a large number of Autonomous Systems (AS) that exchange routing information with each other to learn the best path to the destinations and BGP

(Border Gateway Protocol) [6] is the de-facto inter-AS routing protocol. BGP is a path vector based routing protocol, designed to allow flexible routing policies and be able to adapt to routing dynamics such as link failures and topology changes. Each BGP router establishes a peering session with its neighbor BGP routers, and advertises the entire AS path information to destination prefixes. When a BGP session is set up between two peers, the complete routing tables are exchanged. After this initial exchange, routers only send update messages for additional changes.

Impact of worm attacks on Internet routing infrastructure has been studied before. Researchers have looked at BGP updates during stressful events such as the Code-Red worm attack. [19] first reported the surge of BGP udpates coincided with the Code-Red and the NIMDA worm attack. According to [12], the worm attack had a big impact on some edge networks, and weaknesses in BGP's design and implementation substantially amplified the impact. They reached this conclusion by classifying BGP updates into several categories and examining the cause of each category. Another study [14] also showed that worm attack affected some edge networks like Department of Defense (DoD) networks.

Route flap damping [9] is designed to suppress unstable routes when they flaps frequently, so to prevent the local instability from spreading further. Damping is viewed by the network operation community as one of the major contributors to keep the Internet stable [16], but it has not been fully studied by the research community. In [10], the authors studied the effect of BGP slow convergence on route flap damping. They showed that due to the path exploration in BGP slow convergence, a single flap at one end of the network is able to trigger route damping at another place in the network. This undesired damping will in turn cause longer BGP convergence time.

6 Summary

This paper examined the surge in BGP updates that coincided with the January 2003 Slammer worm attack. Our analysis illustrates how two small edge Autonomous Systems that announce fewer than 0.25% of BGP routing table entries, contributed over 6% of total update messages during the worm attack, as observed from the Oregon RouteViews monitoring points. We also showed that these two Autonomous Systems generated a large number of updates in different ways. The instability at edge ASes can trigger a large number of AS path changes (DPATH updates) in the global Internet, such as those triggered by AS 18296, and can also trigger a combination of different AS paths and changes in other attributes (SPATH updates), such as those triggered by AS 568. In both cases, BGP allows dynamics belonging to local networks to propagate globally. As a result, a small number of edge networks can potentially cause widespread routing overload to the entire Internet.

Route flap damping is the current BGP mechanism to defend against such undesirable global consequences caused by edge instability. Our analysis of BGP udpate data shows that damping has not been fully deployed even within the Internet core. But simple lack of deployment is not the only problem. Our simulation further reveals that partial deployment of damping not only has limited effect but may also worsen the routing performance under certain topological conditions. Even in the case of full de-

ployment, our simulation results show that its effects can be mixed if the edge generates only a small amount of instability. The results presented here are the first step toward a better understanding of BGP instability and route damping. Overall, adding route flap damping or any feature to the global routing infrastructure results in complex and often unexpected behaviors. It remains a research challenge to design a routing protocol that can scale well in a large, dynamic network and insights obtained from understanding BGP instability under stressful events such as the worm attack can help make the Internet infrastructure more stable and scalable.

References

1. CERT Advisory CA-2003-04, "SQL Slammer," http://www.cert.org/advisories/CA-2003-04.html.
2. David Moore et. al., "The spread of the Sapphire/Slammer worm," http://www.cs.berkeley.edu/ nweaver/sapphire/.
3. Internet Health Report, "Sapphire Worm Attack," http://www.digitaloffense.net/worms/mssql_udp_worm/internet_health.jpg.
4. Tim Griffin, "BGP Impact of SQL Worm," http://www.research.att.com/ griffin/bgp_monitor/sql_worm.html.
5. Avi Freedman, "ISP Security Talk, Nanog 2003," http://www.cs.berkeley.edu/ nweaver/sapphire/.
6. Y. Rekhter and T. Li, "A border gateway protocol (BGP-4)," *Request for Comment (RFC): 1771*, Mar. 1995.
7. CERT Advisory CA-2001-19, ""Code Red" Worm Exploiting Buffer Overflow In IIS Indexing Service DLL," http://www.cert.org/advisories/CA-2001-19.html.
8. CERT Advisory CA-2001-26, ""Nimda Worm"," http://www.cert.org/advisories/CA-2001-26.html.
9. C. Villamizar, R. Chandra, and R. Govindan, "BGP route flap damping," *Request for Comment (RFC): 2439*, Nov. 1998.
10. Z. Mao, R. Govindan, G. Varghese, and R. Katz, "Route flap damping exacerbates internet routing convergence," in *Proceedings of the ACM SIGCOMM*, Pittsburg,PA, Aug. 2002.
11. Univeristy of Oregon, "The Route Views Project," http://www.antc.uoregon.edu/route-views/.
12. L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. Wu, and L. Zhang, "Observation and analysis of BGP behavior under stress," in *Proceedings of the ACM SIGCOMM Internet Measurement Workshop 2002*, Nov. 2002.
13. PC World, "Slammer worm slaps Net down but not out," http://www.pcworld.com/news/article/0,aid,108988,00.asp.
14. X.Zhao, M. Lad, D. Pei, L. Wang, D. Massey, and L. Zhang, "Understanding BGP Behavior through a study of DoD Prefixes," in *DISCEX 2003*, Feb. 2003.
15. C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet routing convergence," in *Proceedings of the ACM SIGCOMM 2000*, August/September 2000.
16. Geoff Huston, "Analyzing the Internet BGP Routing Table," *The Internet Protocol Journal*, March 2001.
17. ssfnet.org, "SSFNET modeling the global internet," http://www.ssfnet.org.
18. B.Premore, "Multi-as topologies from bgp routing tables," http://www.ssfnet.org/Exchange/gallery/asgraph/index.html.
19. J. Cowie, A. Ogielski, B. J. Premore, and Y. Yuan, "Global routing instabilities triggered by Code Red II and Nimda worm attacks," Tech. Rep., Renesys Corporation, Dec 2001.