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. This paper examines the surge in BGP updates that coincide with events such as the recent Internet worm attacks. 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 increase in the number of BGP routing update messages observed globally. 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 during the worm attack as observed at the major monitoring points. 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 a number of critical AS-AS peering links were operating above critical load thresholds during the Slammer attack period, and [2, 4, 5] noted that the Slammer worm attack coincided in time with a

* 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.

large increase in the number of BGP routing update messages as observed globally. In fact, such coincidences between Internet worm attacks and the surges in BGP routing update messages have been observed during previous worm attacks as well, such as during the Code-Red [6] attack in July 2001 and again during the NIMDA [7] attack in September 2001.

In this paper, we analyze the BGP [8] 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. Each update received from peer N for prefix p increases a penalty value associated with $\langle N, p \rangle$ and the administrator can set the penalty value for withdraws, advertisements, and duplicate updates. Once the penalty exceeds a configured 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 configured 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, as we will show later in this paper, 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 previous study in [10] and our simulation. 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 the Slammer worm attack coincided in time with a large increase in the number of BGP routing update messages. To understand the worm's effect on BGP, in this section, we delve deeper into the update behavior during the SQL Slammer worm attack in order to investigate 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. The Slammer worm was released on January 25, 2003; we examine BGP update logs from Jan 22nd 2003 to Jan 29th 2003, a period of 7 days around the worm attack. We also apply the techniques from [12] to remove monitoring artifacts from the logs.

We first note that 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 shows the number of route announcements sent by three of the monitored Autonomous Systems. Early on January 25, the number of update spikes dramatically and tails off by January 26th. Note the surge in updates is not localized to a particular AS and the update counts from all three monitored AS rise in similar fashion. Only three ASes are shown for clarity but similar behavior is seen at all monitored AS. Overall, the results show a large surge in BGP updates coincided with the January 25th release of the Slammer worm.

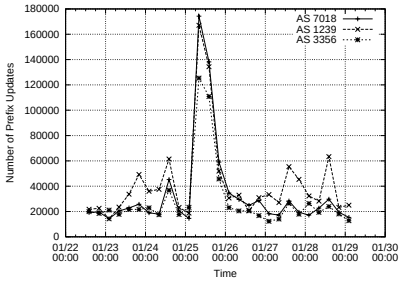

Fig. 1. Path Announcements from Jan 22nd 2003 to Jan 29th 2003

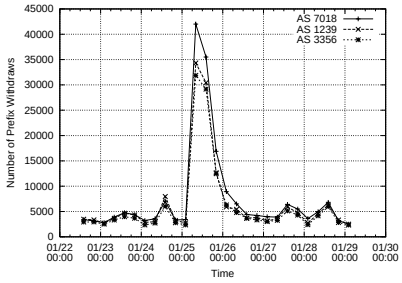

Fig. 2. Withdrawals from three peers from Jan 22nd 2003 to Jan29th 2003

In addition, Figure 2 shows the number of withdraw messages observed during the same worm period. A withdraw message is sent when an AS no longer has *any* route to reach a particular prefix. The dramatic increase in the number of withdraws indicates that, not only did the paths to prefixes change during the worm period, but some prefixes

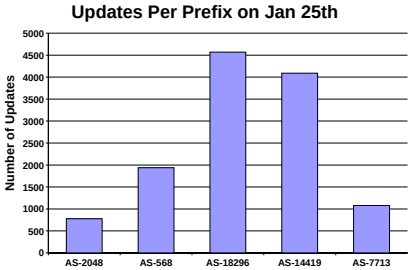

Fig. 3. Average Updates per Prefix - Jan 25th 2003

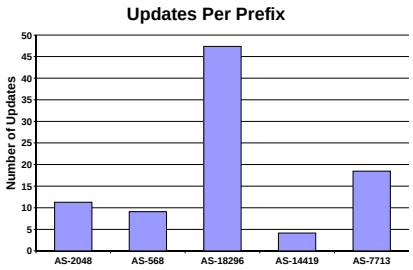

Fig. 4. Average Updates per Prefix - Jan 24th 2003

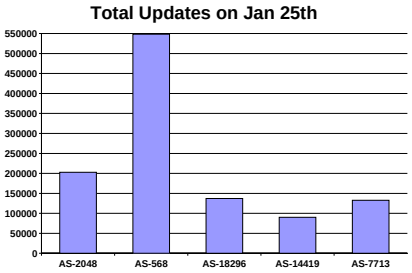

Fig. 5. Total Updates - Jan 25th 2003

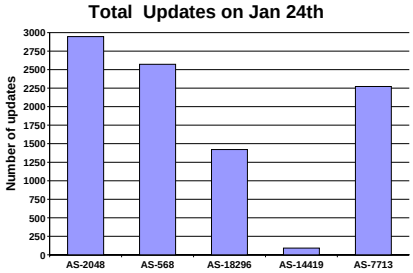

Fig. 6. Total Updates - Jan 24th 2003

were declared unreachable. This suggests that there was a loss of connectivity to some portions of the Internet.

2.2 Identifying Edge Instability

[2] reported that the worm caused high traffic congestion on the edge ASs. Combined with the increase in withdraw messages seen in Figure 2, this suggests that instability at or near the edge AS could trigger at least some of the BGP update surge. 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 25th. Figure 3 shows the average number of update messages per prefix and Figure 5 shows the total updates for prefixes originating from the AS, as observed from all the monitoring points. For comparison, Figure 4 and Figure 6 show the updates per prefix and the total updates on the day prior to the worm attack.

The figures show that prefixes from AS 18296 experienced the highest number of updates per prefix, averaging over 4500 updates per prefix on the day of Slammer attack. However, AS 18296 averaged only 47 updates per prefix on the day prior to the Slammer attack. AS18296 is owned by a university in South Korea and reports in [13] suggested that South Korea's connectivity was among the worst affected by the SQL Slammer worm. AS 18296 advertises only about 30 prefixes and a typical global BGP routing table contains over 120,000 prefixes. 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 Jan25th.

AS568, owned by the US Department of Defense, also stands out in terms of both the average number of updates per prefix and the total number of updates on January 25, 2003. Although AS 568 sends fewer updates per prefix than AS 18296, it advertises more prefixes (a total of 283 prefixes).

Figure 7 shows the percentage of updates contributed by both AS568 and AS18296. This combined set of AS announces less than 0.25% of Internet prefixes, but contributed over 6% of total updates seen during the worm attack.

2.3 Analysis of AS 18296 Edge Instability

To better understand the behavior of AS 18296, Figure 8 classifies the BGP update messages associated with its 30 prefixes as viewed from a particular ISP, AT&T (AS7018). The updates in Figure 8 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.

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 AS 18296 prefixes, indicating a problem near the origin. Furthermore, the AS next to the origin, AS 9318, exhibited few changes in connectivity to destinations expect 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.

2.4 Analysis of AS 568 Edge Instability

AS568 announces 283 prefixes to the global routing table. As viewed from AT&T (AS 7018), AS568 generated less than 100 BGP update messages on the day prior to and a day after Slammer attack, but generated about 27,000 updates on Slammer attack

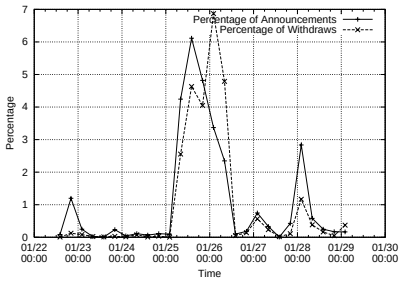

Fig. 7. Percentage of updates on prefixes belonging to AS18926 and AS568 from Jan 22nd to Jan28th

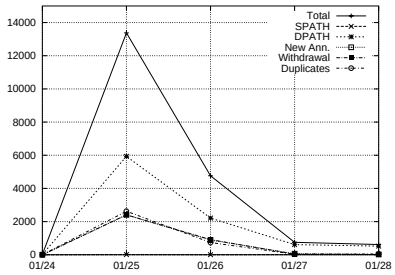

Fig. 8. Classification of Update Messages for AS 18296

day. In this respect, the behavior of AS 568 is similar to that of AS 18296. However, further examination of the update classification shows that, in addition to suffering from unstable routing connectivity, AS568 also experienced additional problems.

Once again we consider the view of AS 568 as seen from a single ISP (AS 7018) and figure 9 shows the type of BGP update messages generated by AS 568's 281 prefixes. Recall that for AS18296, DPATH messages (Different AS Path messages) were the dominant type. But for AS568, 50% of the updates fell into the SPATH (Same AS Path, but difference in some other attribute) category. These SPATH updates signal *no change* in the AS path used to reach the AS568 prefixes.

[14] noted that AS568 also generated an excessively high volume of SPATH updates during the CodeRed/Nimda worm attacks, and most of the updates from AS568 were due to a change in a particular BGP attribute, the AGGREGATOR attribute. Figure 10 shows that once again nearly all of the SPATH updates reported a change in the AGGREGATOR attribute. AS568 has multiple connecting points to its upstream ISP, and the border routers perform internal prefix aggregation before announcing the prefixes to the ISP. The AGGREGATOR attribute contains the value of the border router. [14] attributes this behavior to overloading of some of these links due to the heavy load generated by the worm traffic, ultimately leading to AS568 alternating the links used to reach its provider. This local change should have no impact to any router beyond AS568's ISP. However, because BGP defines AGGREGATOR attribute as a *transitive* attribute and all transitive attributes are propagated globally. As a result, the changes in the local routing connectivity between AS568 and its ISP generated a large number of BGP updates. These updates were propagated to the rest of the Internet, even though the BGP AS path has not changed and the local change would not affect any remote AS beyond AS568's ISP.

The above problem reveals another weak point in the current BGP design. In the case of AS18296, the excessive update messages from a single edge AS were caused by the changes in AS paths to reach the destination. However, In the case of AS568, excessive update messages were propagated globally not only because of AS path unstability, but also due to other local changes. The fact that any single edge AS or small

set of prefixes can cause a spurt of global update imposes a great challenge to BGP's scalability.

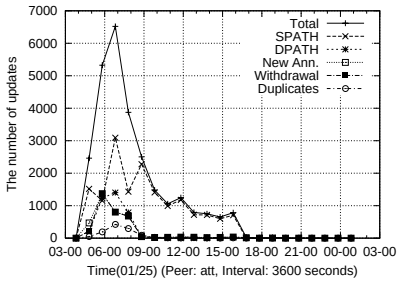

Fig. 9. Classification of Update Messages From AS 568

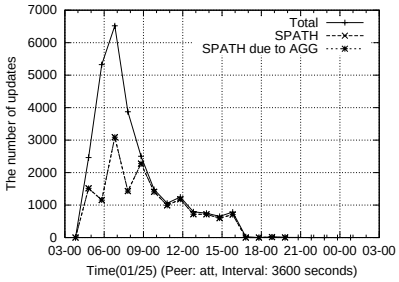

Fig. 10. Aggregator Attribute for Updates From AS568

3 Worm Attack and BGP Damping

In section 2, we observed that the worm caused a high number of updates and identified particular edge AS instability that was partly responsible for the update surge. Concerns over update surges have been raised before, and BGP has features such as MRAI timer and and route flap damping to deal with route instability. In this section, we investigate how well these features fared during the update surge induced by the work attack.

3.1 The MRAI Timer

The Minimum Route Advertisement Interval (MRAI) timer was designed to reduce the number of updates during BGP convergence. After announcing a route to a prefix to its peers, a BGP router will not send any further updates for this prefix until the MRAI timer expires. The intention is that the MRAI timer will suppress routing instability during this hold time, and may help BGP convergence by limiting AS path exploration [15]. The MRAI timer is a standard feature of BGP and is expected to be enabled on most BGP routers. Fig. 11 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, and other messages have larger inter-arrival time. This suggests that MRAI timer was turned on by all neighbors of AS 7018, and likely reduced a large number of updates, which otherwise would have been sent to AS 7018. Similar results were obtained for other monitored AS.

3.2 Route Flap Damping

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

AS Path	Update #	Avg. time of path
7018 3549 9318 18296	120	30.6 s
7018 701 9318 18296	161	228.6 s
7018 9318 18296	7	171.8 s
No Path	93	260 s

Table 1. Path statistics for 203.250.84.0/24

is to suppress the updates caused by an unstable link from propagating further. For each prefix and each peer, a BGP router maintains a penalty value. Whenever a peer makes changes to its route to the prefix, such as changing the ASPATH attribute or withdrawing the path etc., its penalty is increased by a fixed, pre-defined value. Different changes are penalized by different values. The penalty also decays exponentially over time according to the equation

$$p(t) = p(t_0)e^{-\lambda(t-t_0)}$$

The value of λ is often expressed using a *half-life* parameter. When a route's penalty exceeds a *suppression threshold*, a BGP router will stop using this route, thus preventing future changes from propagating further. A suppressed route can be reused only after the penalty drops below a *reuse threshold*. An implementation often has an upper limit on how long a route is suppressed, usually 1 or 2 hours. Table 2 shows the default parameter setting for Cisco routers and Juniper routers.

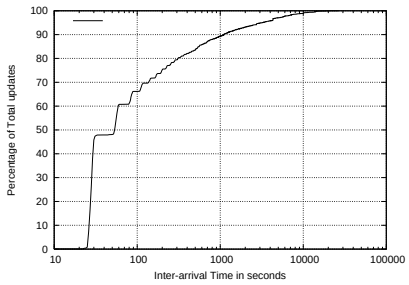

Fig. 11. Cumulative Distribution of Inter-arrival time for AS 18296 for Jan 25th 2003 from AS 7018

[9] recommends the Internet core routers enable damping. 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. However, it is unclear whether damping was effective in dealing with the route instability caused by worm attack.

Damping Parameter	Cisco	Juniper
Withdraw penalty	1000	1000
Readvertisement 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

Table 2. Default parameters for route flap damping (from sigcomm02 paper)

We examined at the behavior of prefixes from AS 18296 (discussed in the previous section). Since all prefixes from AS 18296 behave similarly and generate roughly the same number of updates and withdraws, we show the path statistics for only one prefix, 203.250.84.0/24 and focus on the AS path viewed from AS 7018. Prior to the worm day, this prefix had a steady path of (7018 9318 18296). But during the worm day, this previously stable path was used for only 2% of the time. For the rest of the day, the prefix was unreachable (no path) for 36% of the time, and used path (7018 701 9318 18296) for 55% of the time. Table 1 shows that the AS 7018's path to prefix 203.250.84.0/24 was very unstable, changing once every few hundrends seconds on average. This is beyond MRAI's time scale, and falls in the range where damping should take action. With damping in place, it would be surprising to see such a high number of updates during the worm attack. In the next subsection, we investigate damping during the work attack in more detail.

3.3 Case Study on Damping

We analyzed BGP update data to study how damping might work during the worm attack. All BGP updates we collected were transmitted on the link between the monitoring point and its peering ASes, but this link is not of interest. We want to infer BGP activity on real operational links, which are further away from the monitoring point. Due to the limitations of BGP monitoring and data collection, some assumptions are needed to make the analysis feasible. Like 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.

Case 1 We identifed a sequence of frequent updates received from AS 7018 for the path to AS 18296, and analyzed it to see if route flap damping should have been triggered. Fig. 12 shows the topology containing the relevant ASes and links. The link we are interested is the one connecting AS 7018 and AS 701. We infer BGP updates on this link and calculate its damping penalty over time. If damping is turned on at AS

7018, when the link's penalty is above the suppression threshold, we should not see AS 7018 advertising a path that uses AS 701 until the penalty drops below the reuse threshold. The result presented in this paper is caculated using Cisco's default parameters. A computation with Juniper's default parameters also showed the absence of damping.

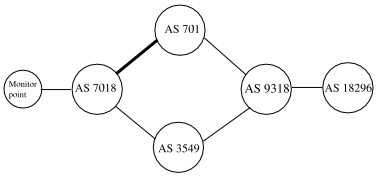

Fig. 12. Topology in Case 1

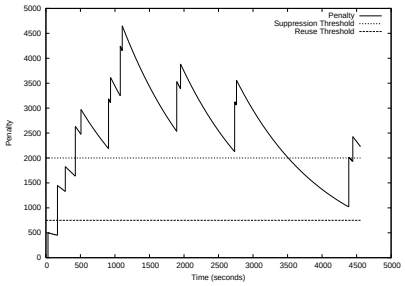

Fig. 13. Damping Penalty on link between AS 7018 and AS 701

Table 3 shows the first few updates in this sequence. It begins with a withdraw, which is not limited by the MRAI timer and should propagate fast in the network. We treat it as an indication of consistent routing state (i.e., no one in Fig. 12 has a path to AS 18296) and start our analysis from this point. The initial penalty value is set to 0. Note any positive values could only make the penalty higher and strenghten our conclusion that damping was not applied. The second update announces a path via AS 701, which can only happen if AS 701 has sent such an announcement to AS 7018 and results in a damping penalty of 500 on the link between AS 7018 and AS 701. 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 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. For example, in this case study, we try to show that damping could have been triggered (i.e., penalty has exceeded the suppression threshold) but actually was not. Therefore, when there are multiple ways to explain an event, we choose the one resulting in the smallest penalty value, so that our conclusion will hold for all the possible choices. For the third and fourth updates in Table 3, the exaplanation we take gives a penalty of 1430.1 at time 167, while the other one has a penalty of 1449.9. The same reasoning leads us to choose an initial penalty value of zero at the beginning of the sequence. Finally, the fifth update is another announcement by AS 701.

Fig. 13 shows the damping penalty value on link between AS 7018 and AS 701 over a time period of about 75 minutes. We can see that the penalty was well above the

No.	Time (second)	Update observed	Possible update on interested link	Induced Penalty	Totoal Penalty
1	0	withdraw	withdraw	-	0
2	30	(7018 701 9318 18296)	announcement	500	500.0
3	141	(7018 3549 9318 18296)	withdraw	1000	1459.0
4	167	withdraw	-	0	1430.1
5	281	(7018 701 9318 18296)	announcement	500	1809.9

Table 3. Sample update sequence and analysis

suppression threshold, reaching as high as 4650 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, we should've seen no update from AS 7018 that 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 as part of the Internet core, AS 7018 did not implement route flap damping at its link to AS 701.

Case 2 Besides cases that suggest lack of damping, we also observed sequences that suggest that damping could have been in effect. In a sequence of seemingly frequent updates for a sample prefix originating from AS 4755 (VSNL), AS 1239's path to AS 4755 keeps flapping, but its damping penalty never exceeds the suppression threshold, and there exists large time intervals in the update sequence that allow the penalty to decay. Fig. 14 shows the toplogy of ASes and links involved. The link between AS 1239 and AS 6762 is the one under study. We start by noting a point where the AS path (1239, 6762, 4755) has been stable for more than 2 hours, thus penalty from previous flaps, if any, would have decayed to a very small number, making the intial penalty of zero a safe assumption. In our analysis, if there are multiple explanations, we choose the one giving the largest penalty value.

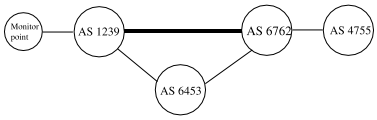

Fig. 14. Topology in Case 2

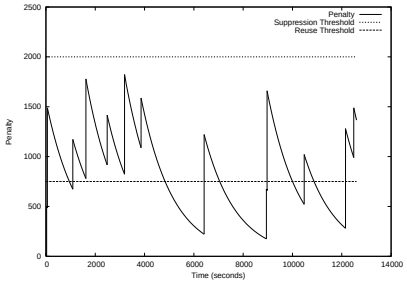

Fig. 15. Damping Penalty on link between AS 1239 and AS 6762

Fig. 15 shows the penalty value of link between AS 1239 and AS 6762 over a time period of more than 200 mintues. During this period, the penalty does not exceed the suppression threshold. Furthermore, detailed examination shows that there are some quiet periods of 20-40 minutes, during which there are no updates at all. Due to the limitation of current BGP monitoring approach, it is hard to infer BGP activity on links further away from monitoring point. However, Fig. 15 and the pattern of update messages suggest that damping might be in effect on link between AS 6762 and AS 4755. When this link is suppressed by AS 6762, both AS 1239 and AS 6453 will not receive any update from AS 4755, which can explain the long quiet periods observed from AS 1239. These quiet periods allow the penalty value to drop significantly before the next penalty is applied, making it never reach the suppression threshold.

4 Effectiveness of Route Damping

The observed BGP data shows that instability at an edge AS can trigger a large number of global updates. Analysis also shows that BGP route flap damping could have an impact on the number of updates, but damping was not fully deployed in the Internet. In this section, we use simulations to evaluate the potential effectiveness of route flap damping in response to instability at an edge AS, similar to the SQL worm attack.

4.1 Simulation Settings

We used the SSFNET BGP simulator [17] to simulate BGP behavior with different topologies. SSFNET was designed to model and simulate the behavior of various network protocols in large networks and includes a standard BGP implementation. In our simulations, besides using standard default values for BGP parameters, 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, which is used widely in real network operations. Note the $MRAI$ timer value plays a major role during convergence, since it is significantly larger than the link delay and processing delay. For route flap damping, we use the default Cisco parameters as described in Table 2.

In this work, we used topologies with 110 nodes and 208 nodes. The topologies are derived from real Internet routing tables, which were processed in a way to approximate the customer-provider or peering relationships of ASes [18]. For clarity, we only present the results for 110 nodes topology. Simulations on the 208 node topology gave qualitatively similar results.

For each simulation run, we first randomly chose one node from the topology and attached an extra AS node to it as the origin AS. 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. 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. We also simulated different rate for the pulses and the results were found to be similar.

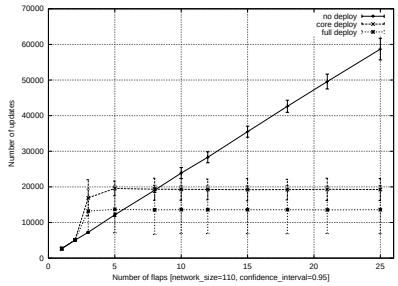

Fig. 16. Number of updates when flapping source is attached to the core

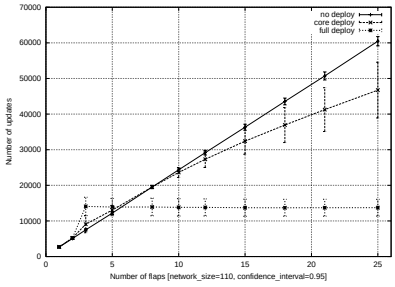

Fig. 17. Number of updates when flapping source is attached to the edge

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 16 and 17 show simulation 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

The BGP route flap damping standard[9] recommends deployment at "core routers". In our simulation, a core node is distinguished by its node degree. We empirically selected nodes with degree larger than or equal to 13 as core nodes, and totally 8 (7.3%) nodes were selected as core nodes. The X-axis represents the number of pulses generated by the flapping source (i.e. the instable edge AS) and for each number of pulse, we run a set of simulations by randomizing the location of flapping source. The results presented in the figures are the average values with 95% confidence interval. The Y-axis represents the total number of updates in the network during the simulation.

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 16 and 17, the number of updates sent with no damping deployed increases linearly as the number of flaps increases. However, with full 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 reduction in total updates can help reduce the routers' load when excessive updates were generated 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 the deployment of 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 the core nodes as the BGP standard recommends, our simulation results show that the 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 16, when the flapping source was attached to core nodes themselves, 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 16. 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. Presently, BGP (Border Gateway Protocol) [8] is the de-facto inter-AS routing protocol. It 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.

In BGP, each BGP router establishes a peering session with its neighbor BGP routers, and advertises the entire AS path information to destination prefixes. After a new BGP session is set up between two peers, the complete routing tables are exchanged between them. After this initial exchange, routers only send update messages for additional changes to its routing table. Thus, the routing table at any state can be constructed by looking at the initial routing table exchanged and the subsequent updates. Update messages can be triggered by various events, including topology change, link failure, policy change and traffic engineering. Some studies have been done on categorizing the types of updates that are received and attributing most updates to a small set of prefixes [19].

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. [20] 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.

BGP has two mechanisms to suppress excessive updates. The MRAI timer limits the rate of sending routing information by buffering the updates for a short time. 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 deployment, 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. 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.
7. CERT Advisory CA-2001-26, “”Nimda Worm”,” http://www.cert.org/advisories/CA-2001-26.html.
8. Y. Rekhter and T. Li, “A border gateway protocol (BGP-4),” *Request for Comment (RFC): 1771*, Mar. 1995.
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. Jennifer Rexford, Jia Wang, Zhen Xiao, and Yin Zhang, “BGP routing stability of popular destinations,” in *Proceedings of the ACM SIGCOMM Internet Measurement Workshop 2002*, Nov. 2002.
20. 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.