# Is BGP Update Storm a Sign of Trouble: Observing the Internet Control and Data Planes During Internet Worms

**Matthew Roughan**
**University of Adelaide**
**matthew.roughan@adelaide.edu.au**

**Jun Li**
**University of Oregon**
**lijun@cs.uoregon.edu**

**Randy Bush**
**IIJ**
**randy@psg.com**

**Zhuoqing Mao**
**University of Michigan**
**zmao@eecs.umich.edu**

**Timothy Griffin**
**University of Cambridge**
**timothy.griffin@cl.cam.ac.uk**

**Keywords:** control plane, data plane, BGP, Internet worms, network measurement

## Abstract

There are considerable reasons to wish to understand the relationship between the Internet's control and data planes in times for stress. For example, the much publicized Internet worms—Code Red, Nimda and SQL Slammer—caused BGP storms, but there has been comparatively little study of whether the storms impacted network performance. In this paper, we study these worm events and see whether the BGP storms observed during the worms actually corresponded to problems in the Internet's data plane. By processing and analyzing two datasets from RIPE, we have found that while BGP update storms occurred in all three worms, the performance of the data plane degraded during the Slammer worm but did not during the Code Red and the Nimda. No direct correlation should be drawn between the degradation of the Internet data plane and the occurrence of a BGP update storm—it may not be a sign of trouble but a sign of the Internet control plane doing its job.

## 1 INTRODUCTION

It is self-evident that problems in the control plane of the Internet will cause data plane performance problems. *Or is it?* Clearly, control plane problems are capable of causing data plane disruptions, but the relationship is not straightforward—for instance, what constitutes a genuine problem in the control plane?

In this paper, we focus on a particular phenomenon—BGP update storms—and study whether or not the occurrence of BGP update storms at the Internet control plane directly maps to the degradation of Internet data plane performance. This study is along the line of researching the relationship of the control plane and data plane, and we view our work in this paper an important study toward this goal.

Note that the Internet health should always be judged by the performance of its data plane, as this is what affects its users. A "storm" of Internet routing updates should not be used as a sign that the Internet is having a trouble. As long as the data plane performance is well, it could actually represent the Internet's proactivity in adapting to large-scale network disruptions to avoid damage to the data plane!

In order to study the data plane performance during such BGP update storms, we study data collected during the spread of three major Internet worms. Although at different level of severity, these worms have all been known to cause significant BGP update storms. They are:

- **Code Red (v2):** The Code Red worm, version 2, started at around 10:00 UTC on July 19, 2001. More than 359,000 computers were infected in less than 14 hours. The spread of the Code Red Worm is described in detail in [1, 2].
- **Nimda:** Just before 12:00 UTC on September 18, 2001, the Nimda Worm began to infect hosts, peaking at around 19:00 UTC at 160,000 unique hosts [3, 4].
- **SQL Slammer/Sapphire:** The most severe of the three worms considered here, the SQL Slammer worm began at slightly before 05:30 UTC on Saturday, January 25 2003, infecting at least 75,000 hosts, with 90% of these infected within 10 minutes [5, 6, 7, 8].

The rest of this paper is organized as follows. Section 2 is on related work. We describe the measurement data we used for our study in Section 3. Section 4 contains some preliminary results. We then introduce a refined methodology in Section 5 to obtain our main results in Section 6. We discuss our results in Section 7 and conclude the paper in Section 8.

This research is funded by NSF award ANI-0221435. Matthew Roughan was partially supported by ARC grant DP0557066. Jun Li is partially supported by NSF award CNS-0520326 and a research grant from Intel Corporation. Point of contact: Dr. Jun Li, lijun@cs.uoregon.edu.

## 2 RELATED WORK

In the case of each of these three worms, detailed analysis has been conducted, both of the worm itself (see citations above), and of the BGP update storm associated with the worm [9, 10, 11, 12, 13, 14, 15]. A very large increase in the average rate of BGP updates was observed to be correlated with the presence of each worm.

In the case of the Slammer worm, there was even a study of network performance during the worm [9]. This study demonstrated that there were performance impacts caused by the Slammer worm, and this is in part responsible for a general belief that the performance impacts during the Code Red, Nimda, and other worms must have been similar (in relation to the BGP update storms observed at these events).

The impact on BGP by events such as electricity blackout or misconfiguration has also been studied [16, 17].

Further, studies have been conducted related to the relationship between the Internet control plane and data plane. Earlier research studied data delivery performance toward a prefix while BGP convergence toward the prefix is happening [18, 19]. Recent work looked at the correlation between BGP instability and path faults [20]. Lately, researchers have also measured the performance of data streams toward a multi-homed sink under routing changes introduced by a BGP beacon [21], and found little correlation between the data plane performance and the volume and duration of BGP updates.

## 3 DATA SETS

The two datasets used in this case study both come from RIPE, `http://www.ripe.net/`. In particular, we use data from their Test Traffic Measurement (TTM) program [22, 23], to obtain data plane performance data for our study, and data from the RIPE Routing Service [24] to obtain BGP update storm information. Both datasets are publicly available. We describe them in detail in Sections 3.1 and 3.2, and in Section 3.3 summarize why they meet our need to study whether or not the whole Internet is in trouble during a BGP update storm.

### 3.1 Performance Measurements

The TTM performance measurements are gathered using special-purpose boxes deployed in a number of different ISPs, predominantly in Europe, but also extending to the USA, and the antipodes (Australia and New Zealand). The measurements have been collected for more than 6 years, and consist primarily of active probe, one-way delay and loss measurements consistent with the IETF's IPPM (Internet Protocol Performance Metrics) Working Group standards, with supplemental traceroute measurements (other measurements such as delay variation are now available but were not during the earlier worm events under study here). The measurement boxes are noteworthy for utilizing GPS to accurately synchronize the clocks to a degree not seen in many Internet performance measurement deployments.

The TTM data used here consists of two periods: the first from June 1st to September 30th 2001 (covering both the Code Red and Nimda worms); and the second from Dec 1st 2002 to February 28th 2003 (covering the Slammer worm). The packet probes used in the measurement were sent using approximate Poisson sampling. At the time of the Code Red and Nimda worms, the probe rate was 90 packets per hour on each path, where the paths formed a clique between around 40 nodes (the number varies slightly over the several months of measurements). At the time of the Slammer worm, the rate was around 120 packets per minute between around 60 nodes. Given the sampling rates, to obtain reasonably precise statistics for the data (in particular the loss rate), we aggregate up to one hour intervals, and consider the mean and variance of the delays, and the mean packet loss rate over these intervals.

The monitors also conduct traceroutes at a rate of approximately 10 per hour, and these are used by the boxes to determine the frequency of route changes affecting the packet probes (that is, we state that a route change occurs if two consecutive traceroutes return different results). Obviously, route changes could occur with finer time granularity and remain unseen by such a measurement process, but the measure reported here provides a lower bound on the amount of rerouting activity during the period of interest. For more details see [22, 23].

### 3.2 Route Measurements

The second dataset used here was obtained from the RIPE Routing Service [24]. This data is similar in nature to the RouteViews data [25], but is obtained from a different set of peers. We use this data here so that we can compare results with earlier studies [11, 12, 9] of BGP updates during worm events. The data has been aggregated to show the rate of updates (total announcements and withdrawals) per minute over all peers. We will typically focus on the rates averaged over one hour periods so as to match the time resolution of the performance measurements described above. We use data over the same time intervals available in the TTM data.

### 3.3 Summary

The two datasets we use are the best publically available datasets that meet our needs to study whether or not the overall Internet data plane is in trouble during a BGP update storm. As we described in Section 3.1, the TTM data provides a good sample for studying whether the Internet data plane is suffering at a given time. Note that we are not interested in studying the possible correlation between the data plane performance of these TTM test boxes and the BGP updates regarding these boxes; instead, we use the performance among TTM test boxes as a sample of the Internet data plane. Furthermore, since our focus is whether the data plane performance will be in trouble at all *when* a BGP update storm occurs—rather than how troublesome the data plane performance will be according to the BGP update storm level, we pay more attention to the occurrence time of a BGP update storm even than the scale of the storm (in this study, essentially it does not matter whether a BGP update storm includes

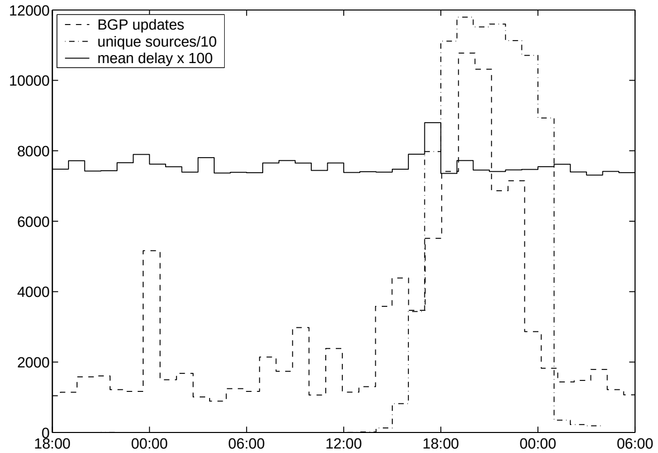

**Figure 1. Three measurements during the peak infection period of Code Red (v2). Note that the three curves are scaled so that they may appear on the same graph.**

or excludes those BGP updates caused by BGP session resets between every pair of a RIPE collector and a BGP router). That is why we carefully ensure that the two datasets are collected by the same organization and that they are consistent with each other, especially in timestamps. These data are also consistent with the data used in previous analysis of the worms in question [11, 12, 9].

## 4 PRELIMINARY RESULTS

Figure 1 shows a comparison between the number of BGP updates, the number of unique hosts actively sending Code Red probes, and the delay in milliseconds of packets on one TTM path. We see a number of important features in this graph. Firstly, it is consistent with studies such as [11, 12], which showed an increase in the number of BGP updates that is highly correlated with the onset of the Code Red worm (the data used here to measure the unique number of sources infected with the Code Red worm was derived from the same set used in [11, 12], namely those reported in [26, 27]). Secondly, we note that the performance on the particular path observed in the TTM data was not significantly affected by the worm.

The two facts by themselves are not very useful. Firstly, one should note that the graphs displayed in Figure 1 are each scaled, so that we may observe them on the same figure, but that the scaling factors are arbitrary. How are we to know whether this scaling is reducing the size of the changes in performance to make them appear less significant, or perhaps, the scaling is making the change in the BGP update process appear more significant than it really is? Secondly, we display performance measurements of only one path of thousands—perhaps this just happens to be one good path, and we are missing the problems. In the following section, we present a methodology to avoid these issues, so we can make genuine direct comparisons between each dataset.

## 5 METHODOLOGY

We wish to perform two tasks: scaling the variations in each dataset so we may make genuine comparisons between these datasets, and providing performance metrics which take into account performance on all of the paths in the network. The following methodology deals cleanly with both issues.

The most important idea to grasp is that of comparison of data to a control set. In medical studies, a control set is used to gauge whether a treatment has a significant impact relative to patients who receive no treatment. Here, we are not applying a treatment, but rather we are seeking to test if the impact of the respective worms was significant. Hence, as our control set, we use the data from periods not significantly impacted by the worm (we choose the entire period excluding the day of a worm and the week following that worm to ensure that no persistent effects of the worm pollute the data; we have verified that even if there are certain worm activities during selected periods, their impact is inconsequential). In this way we can assess how large the observed variations on the day of a worm were with respect to typical variations in the BGP update rate, or the performance along a particular path.

More formally, we perform this process by estimating the mean, and standard deviation of the data (in the control set), and using this to scale the data during the period of the worm. Mathematically, let the delay measurements on path $j$ at time $i$ be given by $X_i^{(j)}$, and let the set $S$ be the set of times considered to be impacted by the worms, which has $|S|$ elements. We use the standard estimators of the mean and variance of the data

$$\bar{X}^{(j)} = \frac{1}{M}\sum_{\substack{i=1\\ i\notin S}}^{N} X_i^{(j)}, \ \ \sigma^2_{X^{(j)}} = \frac{1}{M-1}\sum_{\substack{i=1\\ i\notin S}}^{N}\left(X_i^{(j)} - \bar{X}^{(j)}\right)^2,$$

where $M = N - |S|$, and then we compute the normalized data

$$\tilde{X}_i^{(j)} = \frac{X_i^{(j)} - \bar{X}^{(j)}}{\sigma_{X^{(j)}}},$$

such that the new random variables will now have mean zero, and unit variance (assuming the variance of the data is finite[1]). We perform similar operations on the other measurements: delay variance, loss rate, and number of route changes, and on the BGP data.

The second requirement is to compose the results into a single metric to observe the overall impact of the worms on performance. We do so by considering the averages, for instance for the normalized delay measurements above, i.e.

$$Y_i = \frac{1}{K}\sum_{j=1}^{K}\tilde{X}_i^{(j)},$$

[1] Heavy-tailed distributions (with infinite variance) have often been observed in Internet measurements, however the major impact on these results will be through increased variation of the estimates above, and the introduction of noise into the results as a consequence. A natural way to avoid this issue would be to perform all measurements on the log of the data, but we view the raw data here so as to present results consistent with previous studies.

where there are $K$ paths, so that we have a time series $\{Y_i\}_{i=1}^N$ which also has zero mean. Note that this is, in effect, a weighted mean of the performance data on each path, which gives less weight to paths that are more variable. This key feature of our metric acknowledges that different paths have different degrees of natural variation, and that a direct mean over the data would heavily overweight longer paths (which will generally experience longer delays, and more variability because of the increased number of locations for delay variation), or paths with persistent problems.

Note that while the $\{Y_i\}_{i=1}^N$ have zero mean, they are not guaranteed to be unit variance—the actual variance of these random variables will be dependent on the degree of correlation between measurements on different paths (which is *a priori* known to be non-zero, but its value is not known). Hence, once again, we normalize these random variables to obtain a time series $\{\tilde{Y}_i\}_{i=1}^N$.

From this we can now fix a standard meaning to the variations in $\{\tilde{Y}_i\}_{i=1}^N$. Under assumptions of Gaussian distribution and stationarity, we could predict the exact distribution of $\{\tilde{Y}_i\}_{i=1}^N$, and perform statistical hypothesis tests to determine if a particular measurement (say the delays during the Code Red worm) should be seen as statistically significant variation from the norm. However, the distributions involved are not always Gaussian (we performed tests to demonstrate this, but these are omitted for brevity), and appear to have evidence of non-stationarity (though this is hard to test in the possible presence of long-range dependence [28]), and hence it is difficult to perform precise tests of significance. Instead, we will simply use the 95th percentile confidence bounds (±1.96 standard deviations) to provide a visual indication of the magnitude of the variations with respect to Gaussian assumptions.

## 6 RESULTS

Let us first consider the large-scale pictures shown in Figures 2 and 3. These pictures show the normalized performance metrics (average and standard deviation of delays, loss rate, and number of route changes), aggregated over all paths, on a daily time scale (note that normalization is performed separately on these two periods because of the large separation in time). The important fact to note in Figure 2 is that on the day of the Code Red and Nimda worms (indicated by the vertical dashed lines), there was no noticeable increase in either the delay (both the mean and standard deviation) or the loss rate[2]. This can be determined by observing that the performance curves do not cross the horizontal dotted lines representing ±1.96 standard deviations. More importantly, observing the time series over the large scale, the days of the worms do not appear to be anything special. The number of route changes does cross the ±1.96 lines on both days, but it also does so on other occasions, so we do not know whether to attribute significance to the number of route changes on the days of the worms.

In contrast, the performance data for the Slammer worm in Figure 3 exhibits clear indications of significant increases in the mean delay (on the day of the worm), and the standard deviation of the delay (on subsequent days), and also in the number of route changes occurring on the day. In these cases, once again note that we do not draw significance only from the fact that the performance curve crosses the ±1.96 lines, but also from the fact that the peaks stand out relative to performance on other days.

These two figures give us a reasonable starting point, but they are not adequate for telling the whole story. In looking at the data on large scales, we gain an understanding of its variations over long time scales, but we have obscured the fine details, i.e. it is possible that variations shorter than a day have been obscured. In Figures 4–6, we examine the data at a one-hour time scale over a shorter time interval around each worm. In these plots we restrict our attention to the mean delays for brevity (the other performance metrics do not add much more than can be seen in the large-scale plots). We also directly compare the normalized mean delays to the normalized BGP update rates (note that the absolute values in this graph differ from those in Figures 2 and 3 because the variance of the time series on the hourly and daily time scales is different).

Figures 4–6 clearly show the BGP update increases corresponding to each worm, with the Slammer and Nimda worms generating considerably more activity than the Code Red worm. However, the performance data for the Code Red and Nimda worms show no noticeable performance degradation, but during the Slammer worm, there was a noticeable performance impact that correlates closely to the BGP update activity (consistent with the observed performance impacts reported in [9]).

## 7 DISCUSSION

The take-away result of all of the above is that the performance of the data plane was not well correlated with the BGP update storms observed. In one case, the Slammer worm, the performance degraded in a way directly correlated with the BGP updates, but it did *not* become degraded for either Code Red or Nimda worms, with the latter being particularly significant, because on the same normalized scale, the Nimda worm had a similar magnitude to the Slammer worm. Noticeably, according to the traceroute statistics, the number of route changes affecting the paths of the TTM measurements was clearly significant for the Slammer worm, but only marginally significant for the Code Red and Nimda worms, and perhaps this is an indication of the reason for the performance problems only observed during the Slammer worm. Also, since the Slammer worm was a lightweight, single-UDP-packet attack, it generated far greater congestion than the TCP-based Code Red and Nimda worms, which could also have impacted the performance during the Slammer worm.

As noted earlier, the TTM measurement infrastructure

[2] The peak immediately before the day of Code Red is the result of extremely poor performance measurements on a few paths, relatively briefly, occurring well before any possible start of the worm. We have verified that it was not related to the worm.

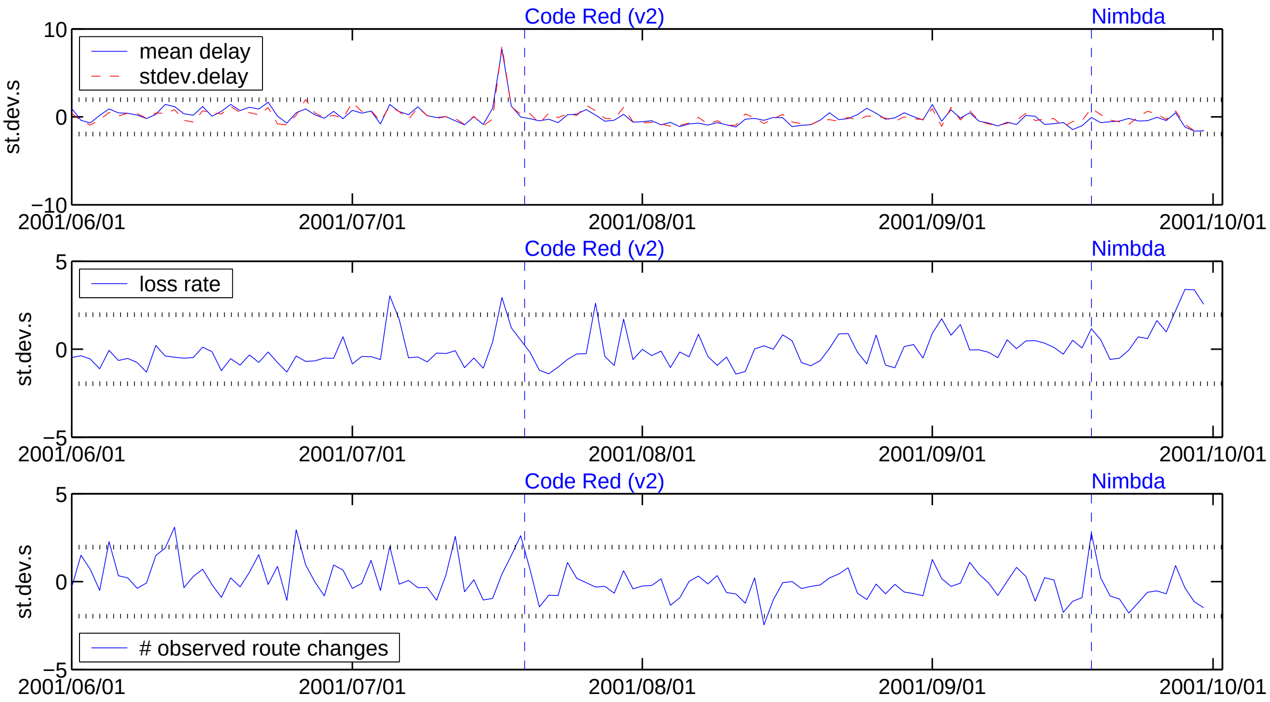

**Figure 2. Three performance metrics derived from the RIPE TTM data for period 1 that includes the Code Red and Nimda worms. The vertical dashed lines show the nominal onset time of the respective worms, and the horizontal dotted lines show** $\pm 1.96$ **standard deviations—the 95th percent significance level.**

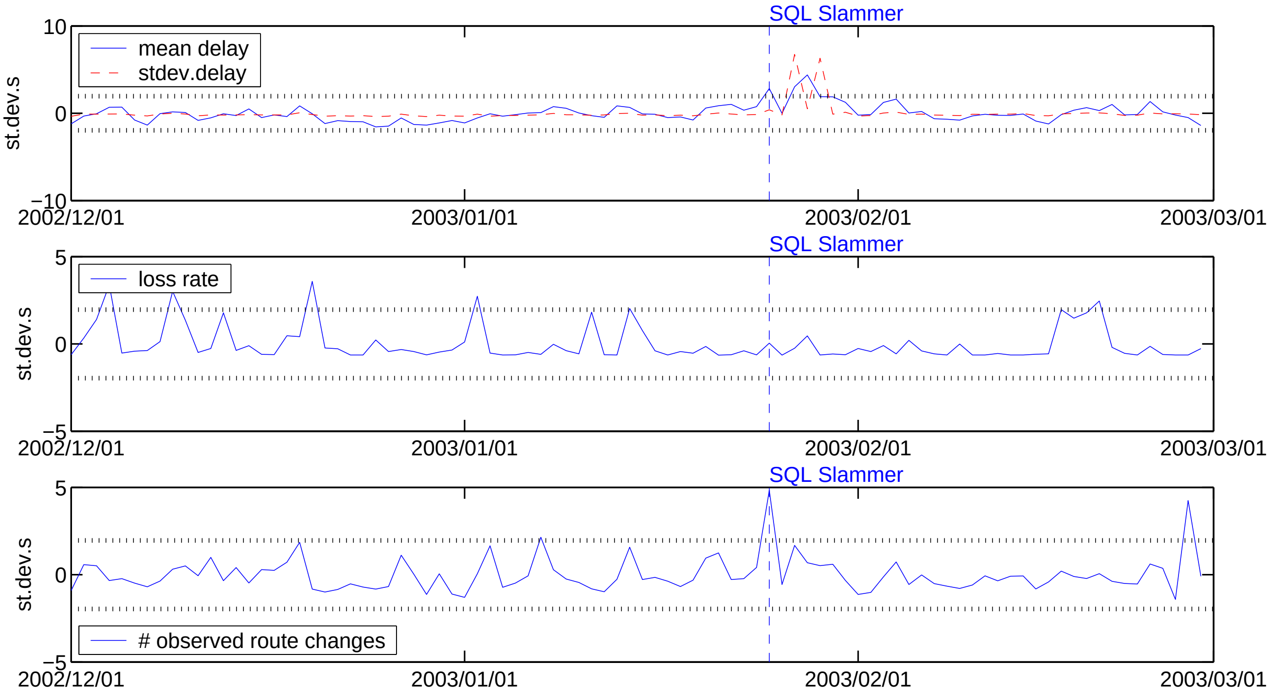

**Figure 3. Three performance metrics derived from the RIPE TTM data for period 2 that includes the Slammer worm. The vertical dashed lines show the nominal onset time of the respective worms, and the horizontal dotted lines show** $\pm 1.96$ **standard deviations—the 95th percent significance level.**

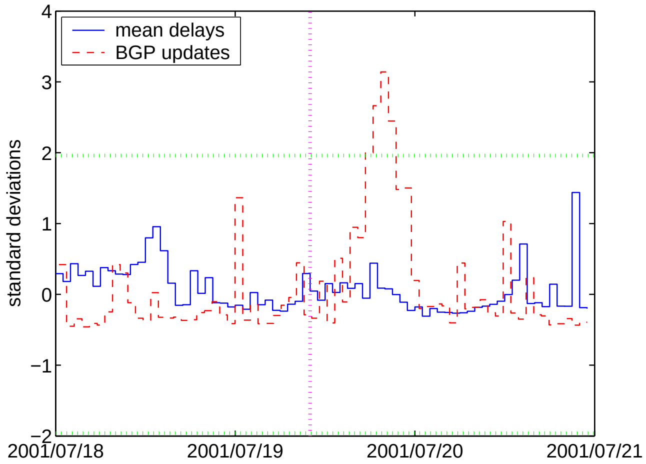

**Figure 4. Results around the Code Red worm, with detailed normalized performance directly compared to the normalized BGP update rate. The vertical dotted lines show the nominal onset time of the Code Red worm, and the horizontal dotted lines show ±1.96 standard deviations.**

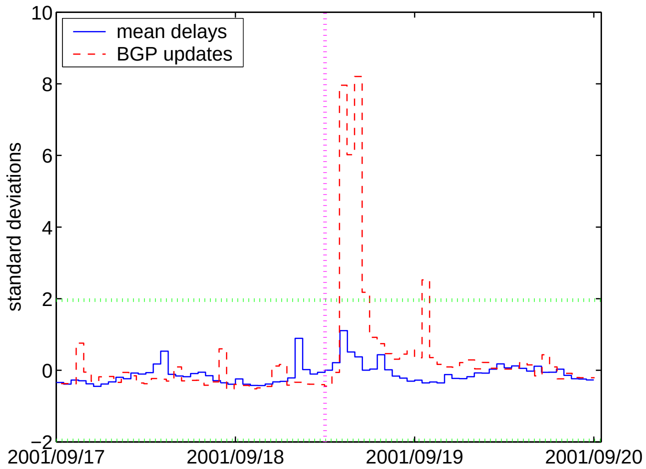

**Figure 5. Results around the Nimda worm, with detailed normalized performance directly compared to the normalized BGP update rate. The vertical dotted lines show the nominal onset time of the Nimda worm, and the horizontal dotted lines show ±1.96 standard deviations.**

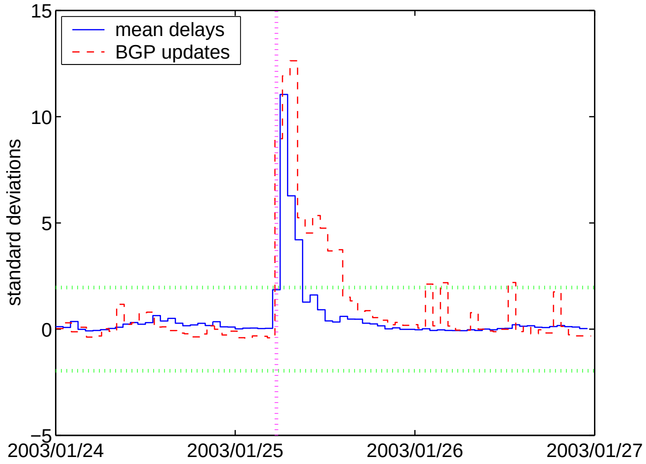

**Figure 6. Results around the Slammer worm, with detailed normalized performance directly compared to the normalized BGP update rate. The vertical dotted lines show the nominal onset time of the Slammer worm, and the horizontal dotted lines show ±1.96 standard deviations.**

grew between the two events, and thus this data source is not perfect for our comparisons. But on the other hand, these changes were not as profound as the difference in the performance measurements; and further, the normalization procedure used here would remove any relative changes introduced (for instance by including more long paths in the data set). It is unlikely that many data collection sources could have remained unchanged over the period of interest (given the disruptions to the global IT economy), but we believe the results used here to be valid within the provisos given.

It is not that BGP updates are unimportant. Clearly their increased presence during these events is of great interest. In at least one case, it was accompanied by simultaneous performance problems. However, we cannot use it as a direct measure of user experiences during the disruptions. A significant increase can be an indication of a severe problem, but more often it may be just an indication that the system is using its built-in controls to correct a problem, before it has an impact on users.

# 8 CONCLUSIONS

Severe global problems, such as Internet worms, could impose intense stress on both the control plane and the data plane of the Internet. In particular, it may cause a sharp increase in the number of BGP updates exchanged between BGP routers, i.e. a BGP update storm.

Since BGP is the routing protocol of the core of the Internet, understanding the implication of BGP update storms is therefore critical. With the data plane performance as the ultimate goal of the Internet, probably the most important question is thus whether or not a BGP update storm would affect the performance of the data plane; and if so, how.

In this paper, we studied BGP update storms during three well-known Internet worms—Code Red, Nimda, and Slammer—and found that while BGP update storms occurred in all three worms, the performance of the data plane degraded during the Slammer worm but did not during the Code Red and Nimda worms. While it is certainly important to pay attention to the occurrence of BGP update storms, our results show that a BGP update storm does not necessarily map to data plane disruption.

Future work includes further investigation on exactly what factors from the control plane caused the data plane degradation during the Slammer worm, especially given that there is no significant degradation during the other two worms. We have also studied the impact on the data plane by artificially introducing routing changes [21], which we call "mild stress," and it would be useful to compare the results from both severe stress and mild stress.

## 9 REFERENCES

[1] D. Moore, C. Shannon, and J. Brown, "Code-Red: A case study on the spread and victims of an Internet worm," in *ACM SIGCOMM Internet Measurement Workshop*, 2002.

[2] CAIDA, "Analysis of Code-Red," http://caida.org/analysis/security/code-red/.

[3] CERT, "CERT advisory CA-2001-26 Nimda worm," http://www.cert.org/advisories/CA-2001-26.html, September 2001.

[4] CAIDA, "Dynamic graphs of the Nimda worm," http://www.caida.org/dynamic/analysis/security/nimda/.

[5] CERT, "CERT advisory CA-2003-04 MS-SQL server worm," http://www.cert.org/advisories/CA-2003-04.html, January 2003.

[6] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver, "The spread of the Sapphire/Slammer worm," http://www.cs.berkeley.edu/~nweaver/sapphire/.

[7] ——, "Inside the Slammer worm," *IEEE Security and Privacy*, vol. 1, no. 4, pp. 33–39, 2003.

[8] "Analysis of the Sapphire worm - A joint effort of CAIDA, ICSI, Silicon Defense, UC Berkeley EECS and UC San Diego CSE," http://www.caida.org/analysis/security/sapphire/.

[9] R. Wilhelm, "TTM and SQL Slammer, impact of the worm attack," in *RIPE 44 Meeting*, January 2003.

[10] J. Cowie, A. Ogielski, B. Premore, and Y. Yuan, "Internet worms and global routing instabilities," in *SPIE 2002*, July 2002.

[11] ——, "Global routing instabilities triggered by Code Red and Nimda worm attacks," Renesys Corporation, December 2001.

[12] J. Cowie and A. Ogielski, "Global routing instabilities," in *North American Network Operators Group (NANOG) 23*, October 2001, http://www.nanog.org/mtg-0110/global.html.

[13] D. McGrath, "Passive Internet health monitoring with BGP," in *North American Network Operators Group (NANOG) 23*, October 2003.

[14] M. Lad, X. Zhao, B. Zhang, D. Massey, and L. Zhang, "An analysis of BGP update burst during Slammer worm attack," in *Proceedings of 5th International Workshop on Distributed Computing*, December 2003.

[15] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. F. Wu, and L. Zhang, "Observation and analysis of BGP behavior under stress," in *Proc. ACM Internet Measurement Workshop*, 2002.

[16] J. Cowie, A. T. Ogielski, B. Premore, E. Smith, and T. Underwood, "Impact of the 2003 blackouts on Internet communications," available at http://www.renesys.com/news/2003-11-21/Renesys_BlackoutReport.pdf, 2003.

[17] S. Misel, "Wow, AS7007!" http://www.merit.edu/mail.archives/nanog/1997-04/msg00340.html.

[18] C. Labovitz, A. Ahuja, A. Abose, and F. Jahanian, "Delayed Internet routing convergence," in *Proc. ACM SIGCOMM*, 2000.

[19] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet routing convergence," *IEEE/ACM Trans. Networking*, vol. 9, no. 3, pp. 293–306, 2001.

[20] N. Feamster, D. G. Andersen, H. Balakrishnan, and M. F. Kaashoek, "Measuring the effects of Internet path faults on reactive routing," in *ACM SIGMETRICS*, 2003.

[21] J. Li, R. Bush, Z. Mao, T. Griffin, M. Roughan, D. Stutzbach, and E. Purpus, "Watching data streams toward a multi-homed sink under routing changes introduced by a BGP beacon," in *Proceedings of Passive and Active Measurement Conference*, Adelaide, Australia, March 2006.

[22] F. Georgatos, F. Gruber, D. Karrenberg, M. Santcroos, H. Uijterwaal, and R. Wilhelm, "Providing active measurements as a regular service for ISPs," in *PAM*, April 2001.

[23] M. Alves, L. Corsello, D. Karrenberg, C. Ogut, M. Santcroos, R. Sojka, H. Uijterwaal, and R. Wilhelm, "New measurements with the RIPE NCC test traffic measurements setup," in *PAM*, March 2002.

[24] RIPE NCC, "RIPE routing information service raw data," http://data.ris.ripe.net/.

[25] "University of Oregon Route Views Project," http://antc.uoregon.edu/route-views/.

[26] D. Goldsmith, Incidents - Internet security incidents mailing list: http://lists.jammed.com/incidents/2001/07/

0149.html, July 20 2001.

[27] K. Eichman, Incidents - Internet security incidents mailing list: http://lists.jammed.com/incidents/2001/07/0159.html, July 20 2001.

[28] M. Roughan and D. Veitch, "Measuring long-range dependence under changing traffic conditions," in *IEEE INFOCOM'99*. NY, NY: IEEE Computer Society Press, Los Alamitos, California, March 1999.