

Correcting Congestion-Based Error in Network Telescope's Observations of Worm Dynamics*

Songjie Wei
Computer & Information Sciences Dept.
University of Delaware
Newark, DE 19716
weis@cis.udel.edu

Jelena Mirkovic
USC Information Sciences Institute
4676 Admiralty Way ste 1001
Marina Del Rey, CA 90292
sunshine@isi.edu

ABSTRACT

Network telescopes have been invaluable for collecting information about dynamics of large-scale worm events. Yet, a telescope's observation may be incomplete due to scan congestion drops, hardware limitations, filtering and presence of NATs, a worm's non-uniform scanning strategy or its short life. We investigate inaccuracies in telescope observations that arise from worm-induced congestion drops of worm scans and show that they may lead to significant underestimates of the number of infectees and their scanning rate. We propose a method to infer worm-induced congestion drops from telescope's observations and use them to accurately estimate global worm dynamics. We apply our methods to CAIDA telescope's observations of Witty worm's spread, and release corrected statistics of worm dynamics for public use.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General;
C.2.3 [Computer-Communication Networks]: Network Operations; C.2.6 [Computer-Communication Networks]: Internetworking

General Terms

Measurement, Security

Keywords

Worm Spread, Network Telescope, Observation Error

1. INTRODUCTION

Network telescopes help researchers observe Internet-wide security incidents. A telescope monitors all traffic sent to an unused but assigned portion of the IPv4 address space. This traffic is assumed malicious since no legitimate services

*This work has been supported by the National Science Foundation, under the grant number 0708774

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'08, October 20–22, 2008, Vouliagmeni, Greece.
Copyright 2008 ACM 978-1-60558-334-1/08/10 ...\$5.00.

are offered by the monitored IPs. Network telescopes have been used to collect information about worm propagation dynamics [1, 2, 3], and about DDoS attacks that use spoofing [4].

Internet-wide worm spread events are also studied through theoretical modeling and simulation, which are necessary to test approaches for early worm detection or to evaluate collaborative defenses. Models and simulators must have high fidelity in reproducing worm's propagation dynamics. This fidelity is tested by comparing the output from a model or a simulator with worm spread dynamics observed by a telescope. To perform this comparison observed worm dynamics must be somehow projected into the global (Internet-scale) view of the worm spread.

Researchers tend to make two assumptions when inferring global worm dynamics from a telescope's observation: (1) large telescopes observe each infected host with minimal delay, thus the observed and the actual number of infected hosts are identical, (2) the number of scans sent by an infectee can be obtained by multiplying its scans received at the telescope by the telescope's size. But these assumptions are clearly wrong! During a large-scale worm spread worm-induced congestion may occur at many points in the Internet and lead to packet drops. Thus a telescope will not observe the portion of scans sent into the Internet corresponding to its size, but potentially a much smaller number of scans that survive congestion. Large congestion loss can also cause unbounded delays at observing infectees that scan at a low rate. Errors in inferring global worm dynamics from telescope's observations propagate into worm models and simulators, which are calibrated to match incorrect data. We propose a method to correct congestion-induced observation errors thus ultimately improving the accuracy of worm models and simulators.

There are other possible sources of observation error in addition to congestion, such as presence of NATs, non-uniform scanning, short lifetime of infectees, network administrator actions or measurement equipment errors. While all these behaviors have been observed in the Internet, we are yet unclear if a telescope's observation can be corrected for them. This is why we focus on correcting congestion-induced errors only in this paper. We validate the accuracy of our inference by correcting CAIDA telescope's observation on Witty worm [5] and comparing our result with the ground truth obtained from forensic analysis [6]. We release detailed estimates of global Witty dynamics at <http://www.cis.udel.edu/~weis/Witty> for public use.

2. RELATED WORK

Weaver et al [7] explore techniques to scale down worm spread events for reconstruction in a limited-size network, e.g., a testbed. It is unclear if this approach can be reversed to scale up a telescope’s observation and infer global worm spread dynamics.

Kumar et al [6] reconstruct the random number generator (RNG) states of each infectee by analyzing the Witty worm’s binary code and the CAIDA telescope’s observation of Witty’s scans. This forensics yields information about each individual infectee: its sending rate, number of disks, duration of the uptime, etc. Hamadeh and Kesidis [8] generalize worm forensics using RNG information. While worm forensics produces more accurate information about worm spread dynamics than our approach, it has a few deficiencies: (1) current forensics techniques work only for worms that use linear congruential RNG, (2) worm binary or source code must be available for dissection, which delays analysis and requires human involvement, (3) results may be inaccurate if the telescope experiences significant packet loss or temporary failure. Our approach complements worm forensics by quickly and automatically inferring global worm dynamics from a telescope’s observation *as the worm spreads* and in situations that do not meet requirements for worm forensics.

Zou et al [9] correct the bias a small telescope may experience in observing number of infected hosts, which stems from the delay between a host’s infection and any of its scans hitting the telescope’s address space. This bias is strictly probability-based and orthogonal to observation errors from scan drops, which we correct in this paper. For telescopes of size $1/8$ probability-based bias should be minor.

3. TELESCOPE OBSERVATION ERROR

We use the term *global observation* to refer to worm spread features — number of infected hosts, scans sent into the Internet and infectees’ scanning rates — that would be observed if we could reliably monitor all Internet hosts. *Local observation* refers to the same features observed at a network telescope.

Precision of a telescope’s observation depends on many factors, including its size and location [10]. The larger a telescope’s size, the better its precision. For example, a telescope covering $1/256$ -th portion of the IPv4 address space should expect to see 1 out of every 256 worm scans, assuming uniformly random scanning and no scan loss. The expected maximum delay to observe an infected host is after it sends 255 scans, which is a small detection delay for fast-scanning worms. In absence of packet drops, a locally observed number of infected hosts at this telescope should be nearly identical to the globally observed one, and a locally observed number of scans (from each infectee and cumulative) should be $1/256$ -th portion of the globally observed one.

The **standard inference approach** [1, 2, 3] is to multiply locally observed number of scans by the telescope’s size to obtain its global observation, and to assume that the locally and globally observed numbers of infected hosts are equal.

A telescope’s precision also depends on the distribution of its size across the IPv4 address space, with distributed telescopes having higher precision and smaller detection de-

lay than continuous ones [11]. Popular Internet telescopes like the ones operated by CAIDA [12] and the University of Wisconsin’s WAIL lab [13] monitor a continuous $1/256$ -th portion of IPv4 space.

3.1 Telescope Error

There are several factors that can introduce error in a telescope’s local observation and make the standard inference approach inaccurate.

Scan drops due to congestion. A telescope may have a large delay in observing an infectee if its scans are heavily dropped due to congestion. Congestion drops also lead to errors in estimating the total number of scans in the Internet. A telescope of size $1/N$ of IPv4 space expects to see $1/N$ portion of all scans. In presence of packet drops this expectation is clearly violated because the telescope cannot see $1/N$ of all scans *sent* by infected hosts — some were dropped.

Figure 1 from [1] shows the number of scans per second received at the WAIL’s $1/8$ telescope during Slammer worm propagation, multiplied by 256. This number peaks at 80 million packets per second. When multiplied by Slammer’s scan size of 404 bytes and divided by 256, the total volume of scans received by WAIL is 1 Gbps, which is same as a widely popular network link rate. It is very unlikely that Slammer infectees were scanning at exactly that cumulative speed. Rather, they were scanning at a higher speed and the telescope’s incoming link limited its observation.

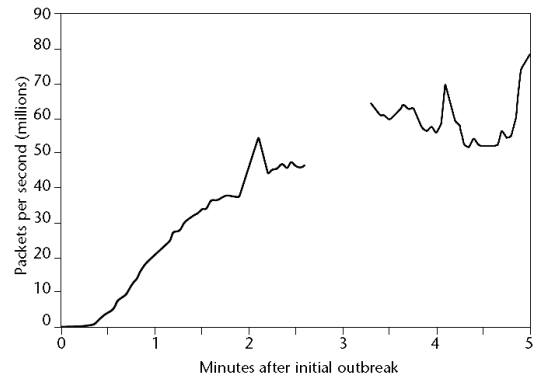


Figure 1: Slammer, observed by WAIL telescope

Can a telescope observe $1/N$ portion of scans *successfully reaching their destinations*? This is highly unlikely either, because congestion on the routes from infectees to the telescope may differ significantly from congestion on the routes from infectees to other destinations. To illustrate this imagine a simplified Internet where core links have limitless bandwidth and all drops occur on the last-hop link to the destination. If the last-hop link to the telescope is 10 times smaller than all other last-hop links, the telescope will see $\frac{1}{10N-9}$ of all received scans¹ — 10 times less than expected!

¹Let p be the probability that a scan will reach its destination, and let it linearly depend on the bandwidth of the last-hop link to the destination. Each scan sent into the Internet arrives at the telescope with probability $\frac{p}{10 \cdot N}$ and it arrives at other destinations with probability $p \cdot \frac{N-1}{N}$. Thus the telescope sees $\frac{\frac{1}{10 \cdot N}}{\frac{1}{10 \cdot N} + \frac{N-1}{N}} = \frac{1}{10 \cdot N - 9}$ of all scans.

Presence of NATs. Many infectees may reside behind a single NAT — their scans appear at the telescope to arrive from a single source, thus skewing the local observation of the number of infected hosts [14]. It is possible to disambiguate NATted hosts using IP ID field information in packet headers [15] for certain operating systems.

Non-uniform scanning. A worm that scans in a non-uniform manner will result in a telescope receiving more or less than its share of scans. Non-uniform scanning could be a result of the worm using subnet scanning, deliberately avoiding telescope’s monitoring, or specifically targeting a telescope. Our inference would apply to such worms if we could correctly estimate the scanning bias, e.g., from a worm’s source code.

Limited life of infectees. An infectee can stop scanning before its scans reach the telescope. Some worms abort scanning after a given time (e.g., Code Red [2]), or they interact with the host machine in a destructive manner that can disable it from sending more scans (e.g., Witty [3]). An infectee can also be deactivated by network administrators.

Measurement equipment errors. A telescope’s measurement equipment may be overloaded and incorrectly record some scans. This is visible from Figure 1 where breaks and irregularities in the recorded scans indicate likely equipment failure.

This paper focuses only on estimating and correcting for telescope observation errors due to worm-induced congestion. Errors from other sources may be equally or more significant than congestion errors, but we are currently unsure if they can be estimated from a telescope’s observations.

3.2 Assumptions about Worms and Internet

1. Constant infectee scanning rate: Unless rate variations are built into the worm’s propagation mechanism, which was not observed in worms to date, an infectee will send out scans at the highest possible rate allowed by its CPU speed, memory, network interface speed, operating system, etc. This maximum sending rate varies for different infectees but remains constant for the same infectee. Our inference holds only for worms that scan at a constant or predictable rate.

2. Packet loss due to worm-induced congestion: Worm scans can be lost due to congestion, routing failure or security actions such as filtering. We assume that majority of packet losses are due to worm-induced congestion so that packets sharing the same routing path experience the same congestion and the same packet loss probability. This is usually true during the growth phase of a fast worm’s propagation, because routing changes and security actions have larger delays than it takes the worm to infect all vulnerable hosts.

3. No significant congestion in the early stage of the worm spread: We assume that there exists an *early stage* of the worm spread, when heavy congestion drops are experienced by none or a few infectees. We say that early stage ends at some time T when congestion starts seriously affecting telescope’s observation. Hosts infected before time T form an *early set* and are called *early infectees*, while the rest are *late infectees*.

While T has no physical meaning, since congestion builds up slowly and differently on different paths, being able to distinguish between hosts infected early or late in the worm spread helps us estimate congestion levels by observing send-

ing rate of those early infectees. As long as T is such that the early set contains majority of hosts that were first observed by the telescope without scan drops, our estimates are accurate. Too large a set leads to lower-than-actual scan rate estimates (due to packet drops that are assumed to be absent), which underestimates global worm dynamics. We infer T from a telescope’s observations, and show in our validation that a range of values produces satisfactory inference accuracy.

3.3 Scan Arrival Ratio – Early Infectees

Scan arrival ratio $P_i(t)$ of an infectee i at time t is the percentage of worm scans sent by this infectee to the telescope that is successfully received. Let an infectee i be observed by the telescope during the early stage of worm spread, when packet drops do not significantly alter observations. It sends $S_i(t)$ packets to the telescope each second, and according to our assumption 1 about the constant scanning rate $S_i(t) \approx \bar{S}_i$, where \bar{S}_i is the constant scanning rate of the infectee i to the telescope.

Let $R_i(t)$ be the number of scans from i received at the telescope at time t and let T be the time when the early stage ends. According to our assumption 3 about no significant congestion in the early stage, we have

$$R_i(t) \approx S_i(t) \approx \bar{S}_i \text{ when } t < T \quad (1)$$

In reality $R_i(t)$ varies a little even when $t < T$ due to packet delay, packet reordering, and random scanning, and we average it to calculate \bar{S}_i as

$$\bar{S}_i \approx \sum_{k=1}^T R_i(k)/T \quad (2)$$

Scan arrival ratio of early infectees, $P_i(t)$, is 1 for $t < T$ and $R_i(t)/\bar{S}_i$, for $t \geq T$.

We now explain how to detect the end of the early stage T and identify early infectees. According to the epidemic model, during the initial stage of a uniform-scanning worm’s propagation, there is an exponential increase in the number of infected hosts [9]. Exponential stage ends either when most of the vulnerable hosts are compromised or when congestion increases so much that it slows down the growth.

We calculate the end of the early stage T as the time when the number of infected hosts departs from the exponential model. This implies that the end of the exponential stage is the same as the end of the early stage. This is true if there is severe congestion during the exponential growth, because it must slow down the growth and thus end both stages. If severe congestion occurs after the exponential growth stage or never, we will underestimate the value of T . This is fine, since the accuracy of our inference depends on two conditions: (1) early set is sufficiently large not to introduce bias and (2) majority of early infectees were first observed during a congestion-free period. Both conditions hold when congestion occurs after the exponential stage.

We measure the match between the number of infected hosts and the exponential model by computing the R-squared value [16], which is a statistical measure to show goodness of fit between a model’s prediction and measured values. We calculate T as the time when the R-squared value starts to continuously decrease from 1, which indicates a decisive departure from the model.

3.4 Scan Arrival Ratio – Late Infectees

We now discuss how to estimate the scan arrival ratio for late infectees. According to our assumption 2 about congestion as the main source of scan drops, we expect infectees that share the route to the telescope to experience same congestion and have the same scan arrival ratio (although their scan sending rates may differ). Thus a late infectee's scan arrival ratio can be obtained from the ratio of an early infectee with whom it shares the route to the telescope.

We assume that infectees that belong to the same BGP atom share the full route to the monitor. This may not be true for large atoms and paths that traverse large ASes, but is the best assumption we can make in absence of detailed, router-level, Internet routing maps, which are not available today. Information about BGP atoms and AS-level routing paths can be obtained from RouteViews [17]. In addition to BGP atoms, we investigated different strategies for identifying infectees which share a route, including grouping them by network prefix. None of these resulted in sufficient amount of route sharing, and thus reduced the accuracy of our inference.

For each late infectee j we attempt to identify an early infectee i within the same BGP atom. On success, we assume that i and j have the identical scan arrival ratio, i.e. $P_j(t) = P_i(t)$. On failure, we calculate the average scan arrival ratio of all early infectees and assign it to the infectee j as its scan arrival ratio:

$$P_j(t) = \frac{\sum_{i \in E} R_i(t)}{\sum_{i \in E} \overline{S}_i} \quad (3)$$

By knowing $P_j(t)$, we can calculate $S_j(t)$ as

$$S_j(t) = \frac{R_j(t)}{P_j(t)} \quad (4)$$

Note that our assumption 1 about the constant scanning rate, \overline{S}_j , still holds, but its estimate, $S_j(t)$, is being recalculated each second. We define the telescope's aggregate scan arrival ratio $\overline{P}_{agg}(t)$ as:

$$\overline{P}_{agg}(t) = \frac{\sum_i R_i(t)}{\sum_i S_i(t)} \quad (5)$$

4. INFERRING GLOBAL OBSERVATIONS

In this section we explore how to accurately estimate global observations of worm dynamics from local ones.

4.1 Number of Infected Hosts

If an infectee i sends $S_i(t)$ scans to the telescope during the t -th second and its scan arrival ratio is $P_i(t)$, the probability that all scans are lost is $(1 - P_i(t))^{S_i(t)}$, which is also the probability that this infectee is not seen by the telescope during the t -th second. In general, this probability should be quite small for infectees that scan at a high rate, and if the congestion is not extremely severe.

Let $Ir(t)$ be the number of infected hosts by the end of time t , and $Io(t)$ be the local observation of this number by the telescope. We define the following variables at time t , observing time as a discrete variable:

$$\Delta Ir(t) = Ir(t) - Ir(t-1) \quad \text{newly infected hosts} \quad (6)$$

$$\Delta Io(t) = Io(t) - Io(t-1) \quad \text{newly observed infectees} \quad (7)$$

$$Uo(t) = Ir(t) - Io(t) \quad \text{unobserved infectees} \quad (8)$$

Due to packet loss $\Delta Io(t) \not\subset \Delta Ir(t)$, but $Io(t) \subset Ir(t)$. Let $\overline{S}_{med}(t-1)$ be the median scanning rate of infectees seen before time t . We could use the average value instead of the median, but this approach favors high-rate scanners, while most of the infectees missed by the telescope are low-rate scanners. We estimate $\Delta Ir(t)$ and $Uo(t)$ as:

$$\Delta Ir(t) = \frac{\Delta Io(t)}{(1 - (1 - \overline{P}_{agg}(t-1))^{\overline{S}_{med}(t-1)})} - Uo(t-1) \quad (9)$$

$$Uo(t) = \Delta Ir(t) + Uo(t-1) - \Delta Io(t) \quad (10)$$

where $1 - (1 - \overline{P}_{agg}(t-1))^{\overline{S}_{med}(t-1)}$ is the probability that at least one scan is observed at the telescope from an infected host. At time zero $\Delta Ir(0) = \Delta Io(0)$, and $\overline{P}_{agg}(0) = 1$. The global observation of the number of infected hosts is

$$Ir(t) = \sum_{k=0}^t \Delta Ir(k) \quad (11)$$

Table 1 illustrates our inference of the global number of infected hosts from local observations. For space reasons values in Table 1 are rounded and we omitted index t from column names. The shaded area contains values that we either set by default (darker shade) or obtain from local observations (lighter shade), while the white area contains values inferred using Eq. (6-11). Column Io is calculated by adding values in the column ΔIo , up to and including time t . For each row (i.e. each value of t) we then calculate in the following order: ΔIr using Eq. (9), Ir using Eq. (11), and $Uo = Ir - Io$.

t	ΔIo	\overline{P}_{agg}	\overline{S}_{med}	Io	ΔIr	Uo	Ir
0	0	1	0	0	0	0	0
1	1	0.8	5	1	1	0	1
2	3	0.5	5	4	3	$9.6 * 10^{-4}$	4
3	2	0.5	5	6	2.06	0.065	6.06
4	3	0.2	5	9	3.03	0.097	9.10
5	6	0.2	5	15	8.83	2.92	17.92

Table 1: Example of inferring Ir

4.2 Number of Scans in the Internet

Let a telescope cover $1/N$ of the IPv4 address space. We can estimate the global scanning rate at time t as the sum of the scans sent by all globally observed infectees at time t . We obtain this by calculating the sum of scans sent by all locally observed infectees at time t , and then use the inferred number of globally observed infected hosts $Ir(t)$ to correct this sum. The global worm scanning rate $C(t)$ is:

$$\begin{aligned} C(t) &\approx \frac{Ir(t)}{Io(t)} \cdot \sum_{i \in Io(t)} S_i(t) \cdot N \\ &= \frac{Ir(t)}{Io(t)} \cdot \sum_{i \in Io(t)} \frac{R_i(t)}{P_i(t)} \cdot N \end{aligned} \quad (12)$$

4.3 Infectees' Scanning Rate

Let a telescope cover $1/N$ of the IPv4 address space. Because of our assumption that infectees send scans into the Internet at a constant rate, we use the maximum of $S_i(t)$

values for estimation of an infectee’s scanning rate B_i :

$$\begin{aligned} B_i &\approx \max_{t \geq 0} (S_i(t)) \cdot N \\ &\approx \max_{t \geq 0} \left(\frac{R_i(t)}{P_i(t)} \right) \cdot N \end{aligned} \quad (13)$$

5. VALIDATION

We validate our method by using the trace of the Witty worm spread as observed by the CAIDA telescope [5], and obtained from the DATCAT repository [18] to infer global observations of number of infectees, number of scans and infectees’ scan rates. We then compare the infectees’ scan rates to the ground truth shown in Kumar et al [6] and inferred via their forensic approach. We are not aware of any other forensic work that would provide additional ground truth for our validation.

The CAIDA telescope monitors a $/8$ block of IPv4 address space. The trace we obtained contains all the worm scans sent to the monitored address space beginning on March 19, 2004 at 4:45 am UTC. All IP addresses in the trace are shown in the original, non-anonymized version. To be comparable with the forensic analysis in [6], we only focus on the first 75 minutes of the trace, including 45.46 million UDP packets with source port 4,000. We apply our inference on infectees that contribute more than 20 scans, as is done in [6].

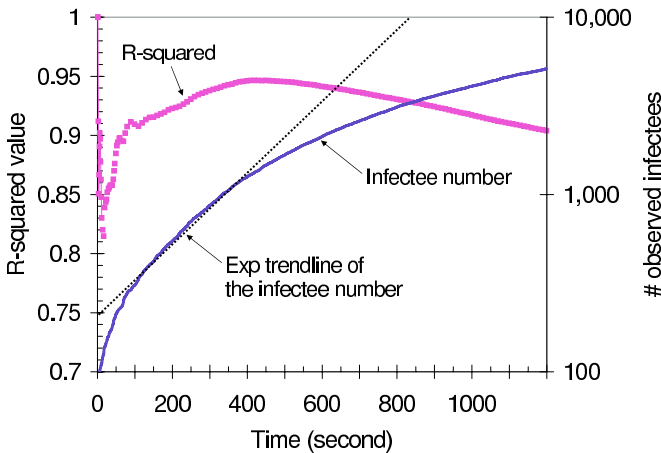


Figure 2: Number of infected hosts

We first identify the early set of infectees by measuring how the telescope’s observation matches the prediction of worm propagation by the epidemic model. Figure 2 shows the evolution of the R-squared value between the observed number of infectees and the derived model, over time. The match quality deteriorates at the very onset of the worm propagation (drop to around 0.8 during first few seconds) because the worm used a hitlist [3, 6]. IPs on the hitlist got infected quickly causing faster-than-exponential growth. Afterwards, the match between the observed and the predicted number of infectees improves, which is shown in the increase of the R-squared value. The prediction starts to deteriorate again after the first 400 seconds, and falls continuously, indicating the end of the worm’s early (exponential growth) propagation stage. Infectees observed before 400 seconds form our early infectee set.

We now compare our estimated distribution of Witty infectees’ access bandwidth with the ground truth obtained from the forensic analysis in [6]. We first estimate the maximum scan sending rate of each infectee (Section 4.3) and then multiply this rate by the average Witty packet size (1,070 B) to obtain access bandwidth.

Figure 3(a) shows the ground truth, our inference of the infectee’s access bandwidth for several values of T , and its standard inference. Standard inference underestimates infectees’ bandwidths by 2-3 orders of magnitude. Our estimate with $T = 400$ is much closer to the ground truth, but it lacks “steps” that are evident there. The reason for this lies in our estimate of the scan arrival ratio of late infectees. If many late infectees cannot be paired with an early set infectee (there is little route sharing), their individual differences are lost by approximating their scan arrival ratio with the average from the early set.

Another error in our estimate occurs for those slow scanners with small access bandwidths. If the telescope only receives scans from an infectee occasionally and lacks information for most of the observation periods, our approach to taking the maximum observed scan rate as the basis for calculation of the infectee’s scan sending rate results in an overestimate. Such infectees contribute no more than 3% of the total scans received by CAIDA telescope thus our inference of the overall scanning rate should be correct (Figure 3(b)).

Figure 3(b) shows our estimate of the Internet-wide scanning rate and compares it to the values obtained via the standard inference. The standard inference line flattens after 800 seconds, which is the time when the 100 Mbps incoming link of the CAIDA telescope is saturated. Thus the highest observed rate, scaled up by standard inference, is 3 million scans per second. Our inference with $T = 400$ detects and corrects the effect of this last-hop congestion. At the end of 75 minutes we estimate 10 million scans per second, which is more than a three-fold difference!

To investigate the effect of different early sets on estimation accuracy we show inference with several T values in Figures 3(a) and 3(b). A larger T , such as 2,000, tends to underestimate the packet loss, because many infectees that experience congestion from the start are erroneously placed into the early set. This results in smaller inferred infectee access bandwidths and underestimates scan rates. A smaller T results in estimates that may be slightly larger or smaller than the ground truth (lines for $T=50$ and $T=200$ in Figure 3(a)) but differences are minor. These results indicate that T does not need to be inferred precisely as long as it precedes the onset of severe congestion.

Our estimate of the number of infected hosts in the Witty trace with $T = 400$ is very close to the telescope’s original observation (graph not shown for space reasons). Telescope sees 11,326 hosts and we correct this to 11,516. Observation error is small because the Witty worm was relatively slow, with a small vulnerable population and thus did not create excessive congestion.

6. CONCLUSIONS AND FUTURE WORK

Congestion during fast worm spread distorts a telescope’s observation and leads to an incorrect picture of global worm dynamics, such as underestimate of the number of infected hosts, number of scans and infectees’ scanning rates. This observation error propagates into worm models and simula-




Figure 3: Correction of CAIDA telescope's observation of Witty's spread

tors that attempt to match inferred global worm dynamics.

We proposed an innovative approach to estimate congestion packet loss from a telescope's local observation, and to use this loss value to correctly estimate global worm dynamics. Our validation using CAIDA telescope's Witty worm observations and the ground truth presented in [6] shows that: (1) our inference matches the ground truth, and (2) our inference is significantly more accurate than the standard inference. We hope that these findings will lead to correction of past and future telescope observations, and to better worm models and simulators.

Our future work will investigate how a telescope's size influences our inference accuracy, and if other sources of telescope error such as NATs, filtering, non-uniform scanning and measurement equipment errors can be inferred from a telescope's local observations.

7. REFERENCES

- [1] David Moore, Vern Paxson, Stefan Savage, Collen Shannon, Stuart Staniford, and Nicholas Weaver. Inside the Slammer Worm. *IEEE Security and Privacy*, Jul/Aug 2003.
- [2] David Moore, Collen Shannon, and kc claffy. Code-Red: a Case Study on the Spread and Victims of an Internet Worm. In *Proc. the Second Internet Measurement Workshop (IMW)*, Nov 2002.
- [3] Colleen Shannon and David Moore. The Spread of the Witty Worm. *IEEE Security and Privacy*, 2(4):46–50, Aug 2004.
- [4] David Moore, Geoffrey Voelker, and Stefan Savage. Inferring Internet Denial of Service Activity. In *Proc. of USENIX Security Symposium*, Aug 2001.
- [5] Colleen Shannon and David Moore. The CAIDA Dataset on the Witty Worm. http://www.caida.org/data/passive/witty_worm_dataset.xml/.
- [6] Abhishek Kumar, Vern Paxson, and Nicholas Weaver. Exploiting Underlying Structure for Detailed Reconstruction of an Internet-scale Event. In *Proc. of ACM Internet Measurement Conference*, Oct 2005.
- [7] Nicholas Weaver, Ihab Hamadeh, George Kesidis, and Vern Paxson. Preliminary Results of Using Scale-Down to Explore Worm Dynamics. In *Proc. of ACM CCS Workshop on Rapid Malcode (WORM)*, Oct 2004.
- [8] Ihab Hamadeh and George Kesidis. Toward a Framework for Forensic Analysis of Scanning Worms. In *Proc. of ETRICS International Conference*, Jun 2006.
- [9] Cliff. Zou, Lixin Gao, Weibo Gong, and Don Towsley. Monitoring and Early Warning for Internet Worms. In *Proc. of ACM CCS*, 2003.
- [10] David Moore, Collen Shannon, Geoffrey Voelker, and Stefan Savage. Network Telescopes: Technical Report. *CAIDA technical report*, 2004.
- [11] Moheeb Abu Rajab, Fabian Monrose, and Andreas Terzis. On the Effectiveness of Distributed Worm Monitoring. In *Proc. of USENIX Security Symposium*, Aug 2005.
- [12] CAIDA. Network Telescope. <http://www.caida.org/research/security/telescope>.
- [13] University of Wisconsin-Madison Advanced Internet Lab. Web page. <http://wail.cs.wisc.edu>.
- [14] Moheeb Abu Rajab, Fabian Monrose, and Andreas Terzis. On the Impact of Dynamic Addressing on Malware Propagation. In *Proc. of ACM CCS Workshop on Rapid Malcode (WORM)*, Nov 2006.
- [15] Steven M. Bellovin. A Technique for Counting NATted Hosts. In *Proc. of Second Internet Measurement Workshop*, Nov 2002.
- [16] George Cassella and Roger L. Berger. *Statistical Inference*. Duxburg Press, 2nd edition, 2001.
- [17] University of Oregon. Route Views Project. <http://www.routeviews.org>.
- [18] CAIDA. Internet Measurement Data Catalog. <http://www.datcat.org>.