Observing Internet Worm and Virus Attacks with a Small Network Telescope

Uli Harder Matt W. Johnson Jeremy T. Bradley
William J. Knottenbelt [1]

Department of Computing, Imperial College London, South Kensington Campus, London SW7 2AZ, United Kingdom

Abstract

A network telescope is a portion of IP address space dedicated to observing inbound internet traffic. The purpose of a network telescope is to detect and log malicious traffic which originates from internet worms and viruses. In this paper, we investigate the statistical properties of observed traffic from a passive Class C telescope over a total of three months. We observe that only a few IP sources and destination ports are responsible for the majority of the traffic. We also demonstrate various ways to visualise the traffic profile from a telescope. We show that specific profiles can identify and distinguish portscans, hostscans and distributed denial-of-service (DDOS) attacks. Looking at the inter-arrival time of packets, the power spectrum and the detrended fluctuation analysis of the observed traffic, we show that there is very little sign of long-range dependence. This is in stark contrast to other network traffic and presents exciting possibilities for identifying malicious traffic purely from its traffic profile.

Key words: Internet worm attack, malware monitoring, network telescope

1 Introduction

In recent years there has been a large growth in internet traffic generated by malware, that is, internet worms and viruses. This traffic usually only impinges on users when, either their machine gets infected, or the internet becomes unusable due to overloaded routers during the epidemic stage of a new worm [1]. However, many users are unaware that there is also a continuous background level of malware traffic at times of non-epidemic growth, so that devices connecting to the internet today are subject to a steady stream of *portscans*, *back-scatter* from attempted distributed denial-of-service attacks

[1] Email: {uh,mwj,jb,wjk}@doc.ic.ac.uk

This is a preliminary version. The final version will be published in Electronic Notes in Theoretical Computer Science
URL: www.elsevier.nl/locate/entcs

and *hostscans*. Recognising and classifying this traffic at an early stage is an important activity if we are to build better firewalls, protect the internet router infrastructure, and provide early-warning mechanisms for new attacks.

In this paper, we present traffic analysis results from an existing network telescope which has been set up to log all inbound traffic that is sent to its portion of internet address space. The telescope we use in this study collects traffic sent to a class C network (256 addresses) that has not been in use for over four years. The use of larger telescopes has been described in [2,3]; however, in this study we show that even on such a relatively small telescope classification of malware-generated traffic is still practical. The monitoring is entirely passive as there are no machines actively connected to the network. Other telescopes make use of so called "honey-pots" to attract traffic; for more details see for instance [2,3,4,5]. In these setups, the telescope interacts with the incoming traffic, allowing more information on the nature of the attacks to be extracted based on the contents of the packets received. In the case of TCP-based attacks, our telescope sees only the start of a conversation, or the first packets returned in case of spoofed IP addresses.

We use `tcpdump` to monitor the incoming data. The data is then stored in a PostgreSQL database to simplify the analysis. Our observations are over two separate time periods: from 28 January 2005 to 5 March 2005 (period 1) and from 24 March 2005 to 13 May 2005 (period 2).

In Section (2) below, we investigate the total traffic and present rank-frequency plots for the destination port and source IP numbers. We also show the inter-arrival time distribution of packets and analyse the aggregated packet rate time series for the long-range dependence that is usually found in network traffic [6]. We also show three-dimensional plots that clearly identify different attack types (similar to [7]).

In Section (4), we focus on attacks that are likely to be caused by the Sasser worm [8]. We try to establish the time between attacks on our class C network. We also show that this particular traffic type has periodicity but no long-range dependence in its packet trace.

In Section (5), we focus on a particular event in the trace that happens in period 1 between 17 and 19 February 2005. We observe a sudden increase in the packet rate where the packets appear to be coming from a webserver.

In the last section, Section (6), we present some simple calculations that show the difficulty of interpreting the network telescope data, by trying to calculate the observed rate for a particular attack from published data.

2 General remarks

In period 1, there were 18,060,643 packets observed, and 20,669,098 in period 2. This yields a mean arrival rate of between 5 and 6 packets per second. We do not observe much variation from this mean packet rate over time, as can be seen in Fig. (2). Tab. (1) shows the distribution of traffic type as observed in periods 1 and 2 by the network telescope: the vast majority of the traffic is TCP, just 5% is UDP and 1% is ICMP traffic. This distribution hardly

Type	Frequency in %
S	90.4
UDP	4.8
NBT	2.0
R	1.6
ICMP	1.2

Table 1
Traffic type distribution of the traffic in period 1 and 2.

differs between the two periods. Similarly, the destination ports (ports of IP addresses inside the telescope) show that only a few ports are used for the majority of the traffic – see Tab. (2). As a consequence, firewalls can be very

Rank	Port # per'd 1	cum. frequency	Port # per'd 2	cum. frequency
1	135	0.39	135	0.42
2	445	0.65	445	0.64
3	1433	0.71	1433	0.74
4	1025	0.75	139	0.78
5	80	0.79	1025	0.83
6	139	0.83	38293	0.87
7	38293	0.86	80	0.91
8	26943	0.90	137	0.93
9	137	0.92	ICMP	0.95
10	6129	0.93	6129	0.95
11	2745	0.94	2745	0.96
12	ICMP	0.95	1434	0.96

Table 2
Cumulative frequencies of destination ports in period 1 and 2.

efficient in keeping out attackers by simply denying a few ports rather than a blanket blackout. We also see that the top three ports responsible for 3/4 of the entire traffic do not change at all from period 1 to 2, and for the top twelve accounting for 96% of the traffic, there is little change in the ranking. The source ports show a small preference for port 80, ICMP and some other ports, which is presumably due to spoofing of the telescope addresses, but the distribution looks much flatter. The large number of port 80 requests are due to backscatter of a possible denial-of-service attack, as we shall see in Section (5).

To get a better picture indicating how differently destination ports are distributed for the two periods we can look at their rank-frequency plot, shown in Fig. (1). This type of plot is often used to illustrate Zipf's law for word frequency [9]. Whilst there is no universal behaviour, there are clearly three

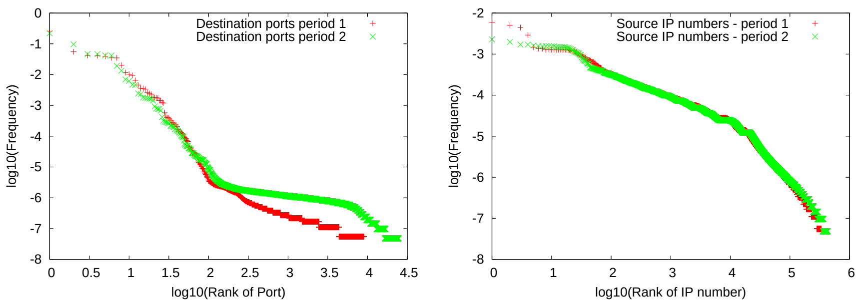

Fig. 1. Rank-frequency plot for the destination ports and the IP numbers in period 1 and 2.

separate regions for the destination ports. The first region is formed by the ten most frequently used ports. After that there are two regions of different power laws. For the source port the distribution is much flatter, which is presumably due to the fact that these are mainly allocated at random by the operating system running on the source host.

A similar picture presents itself when we look at the rank-frequency plot of the source IP addresses in Fig. (1). The source IP addresses show three different regions similar to the destination port numbers.

As we mentioned before, the fact that there are heavy tailed distributions at work can be utilised when one wants to cut down on incoming malware traffic: only a small number of IP addresses and ports are used for the majority of the traffic. However, the downside is that weeding out all worm and virus traffic by listing the IP and port numbers will be hard work as the list is continually growing.

It is also interesting to look at a summary of the total data received by the network telescope during the observation period. To this end, we summarise the data by presenting the average number of packets seen for each hour of the observation in Fig. (2). One can clearly see the traffic intensity change with a frequency of about one day. This might relate to the fact that a lot of the traffic is caused by infected PCs which are switched on and off with that frequency. We can also see an extraordinary peak around the 22nd day of the observation. The traffic intensity suddenly shoots up by an order of magnitude for a brief time. We shall investigate this later.

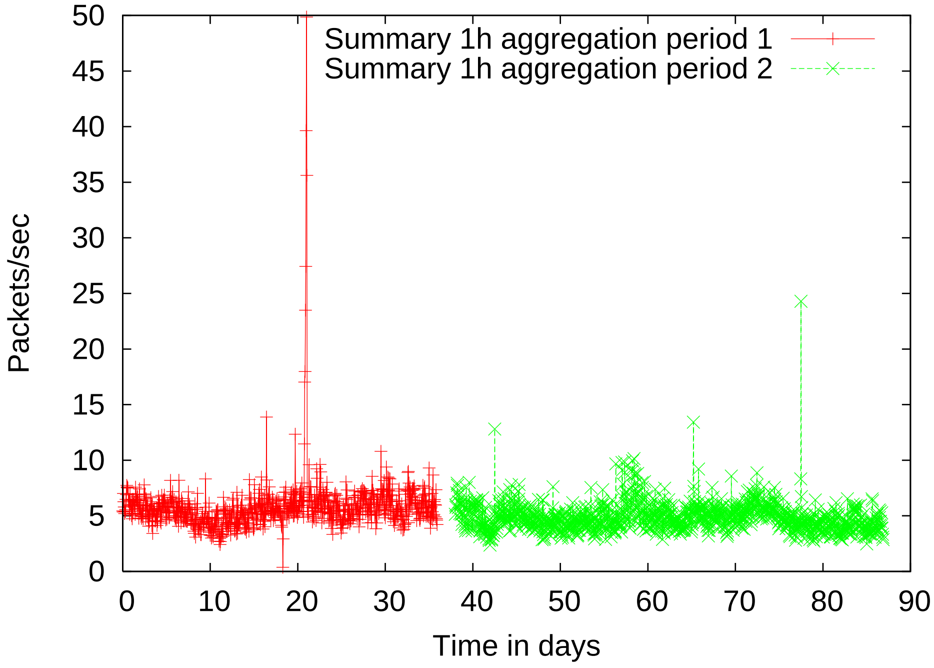

Fig. 2. Summary of the observed network traffic.

In Kim *et al.* [7], it was proposed to use a 3D plot using source IP number, destination IP number and destination port to track attacks. In these plots host and portscans appear as horizontal lines. Denial-of-service (DoS) attacks tend to form squares. We can see in Fig. (3) that our data clearly shows various hostscans having taken place during the monitoring period. In Fig. (3), the

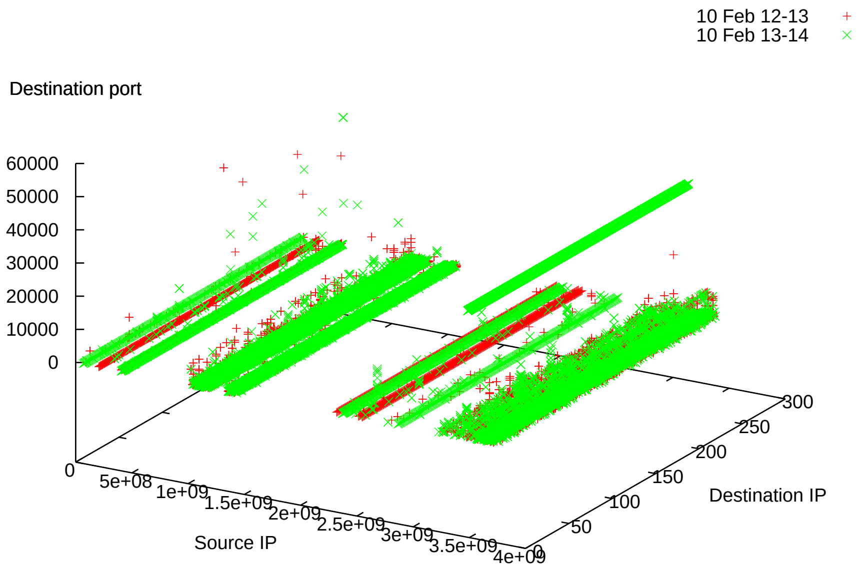

Fig. 3. 3d plot of 2h of observation data.

source IP address is represented as an integer by converting A.B.C.D into $A \times 256^3 + B \times 256^2 + C \times 256^1 + D \times 256^0$.

3 Inter-arrival times and long-range dependence

To investigate the statistical nature of the traffic seen by the network telescope we treat the entries of the `tcpdump` logfile as realisation of a point process. Each packet arrival has an associated timestamp with microsecond accuracy and other properties such as:

- source and destination IP address
- source and destination port number
- type: TCP, UDP, ICMP

We also record the actual content of the packet, but since the vast majority of packets captured are merely TCP SYN or SYN/ACK, and we make no effort to respond to a TCP SYN to bring up a full connection and begin data transfer, this provides us with no useful information for TCP data. In our investigation we look at the entire traffic profile, as well as filtering the data by one or more of the properties above.

We plot the inter-arrival times of packets in a double-logarithmic plot with exponential bin sizes, similar to measurements of network and printer traffic in [10,11]. The overall inter-arrival times show neither signs of being scale-free nor exponential. They do differ from *normal* network traffic though, as there are more extreme inter-arrival times and the distinctive peaks around 10 and 120 microseconds are missing.

Ordinary network traffic usually exhibits long-range dependence. To this end, we investigate aggregate packet counts for bins of 0.1 second size. The most direct way to investigate whether a time series is auto-correlated is to calculate its auto-correlation function (ACF). However, the ACF is very susceptible to trends and other non-stationarities in the data. Better estimates for the long-range dependence are, for instance, power spectrums [12] with filters, or detrended fluctuation analysis (DFA) [13,14,15].

Here we use both the power spectrum and the DFA techniques to analyse the data.

Given a discrete time series x_t, $0 \le t \le N$ with mean $\langle x \rangle$ the auto correlation function at lag s can be defined as

$$C(s) = \langle \bar{x}_i \bar{x}_{i+s} \rangle = \frac{1}{N-s} \sum_{i=0}^{N-s} \bar{x}_i \bar{x}_{i+s}. \tag{1}$$

where $\bar{x}_i = x_i - \langle x \rangle$. For an uncorrelated time series $C(s)$ is zero, for short range correlations the ACF will decay exponentially with a characteristic scale $s_\times$, $C(s) \approx \exp(-s/s_\times)$. In contrast, long range correlations follow a power law in their decay: $C(s) \approx s^{-\gamma}$ with $0 < \gamma < 1$. Using the ACF of Eq. (1) directly usually presents problems as real data will contain trends which, for

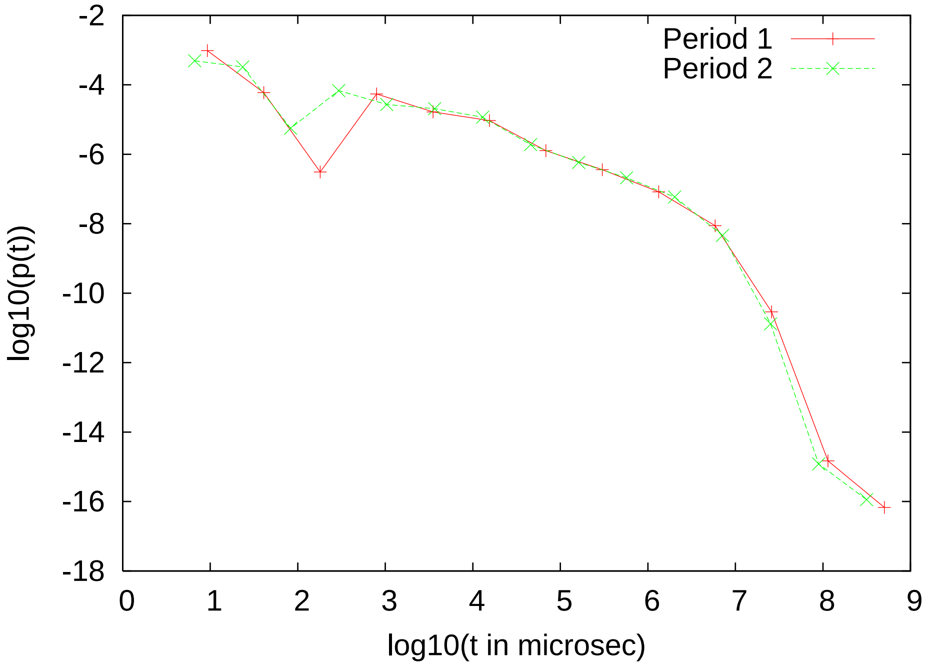

Fig. 4. inter-arrival time distribution.

instance, make the use of the overall mean of the time series in Eq. (1) problematic (more detail on this can be found in [16]).

The fact that long-range correlations tend to imply non-stationarity prohibits the use of the ACF. Other methods, like the power spectrum, also suffer from this problem and are biased estimators. However, the power spectrum can be modified to counter this by the use of filters or windows [12,17].

The detrended fluctuation analysis solves this problem in the following way:

(i) First the profile of the time series is calculated:

$$Y(i) = \sum_{k=1}^{i} x_k - \langle x \rangle \tag{2}$$

(ii) Next we divide the profile into $N_s = N/s$ non-overlapping pieces of length s; for simplicity we assume that there is no remainder ([16] explains how more general situations are treated).

(iii) For each piece $0 \leq \nu \leq N_s$ we now fit a polynomial of degree k with values $p_\nu^k(i)$. This allows the computation of the detrended time series profile for each piece ν:

$$Y_s(i) = Y(i) - p_s^k(i) \tag{3}$$

The degree k of polynomials $p_\nu^k(i)$ can be used to investigate the dependence of the DFA on trends of order k.

(iv) The next step is to calculate the variance of each piece ν:

$$F_s^2(\nu) = \langle Y_s^2(i)\rangle \frac{1}{s}\sum_{i=1}^{s} Y_s^2[(\nu-1)s+i] \tag{4}$$

(v) The final step is take the average of all variances at time scale s:

$$F^k(s) = \left[\frac{1}{2N_s}\sum_{\nu=1}^{2N_s} F_s^2(\nu)\right]^{1/2} \tag{5}$$

where k corresponds to the order of the polynomial.

Data that is long-range power law related then exhibits a power law for the dependence of $F(s)$ on s:

$$F^k(s) \approx s^{\alpha} \tag{6}$$

for large s. As shown in [16] the power law of the ACF and the DFA are related by:

$$\alpha = 1 - \gamma/2 \tag{7}$$

Similarly, the power spectrum of the time series can show a power law

$$S(f) \approx f^{-\delta} \tag{8}$$

where $\delta = 0$ is white noise and $\delta = 2$ Brownian motion. This is related to the DFA in [18] by:

$$\delta = 2\alpha - 1 \tag{9}$$

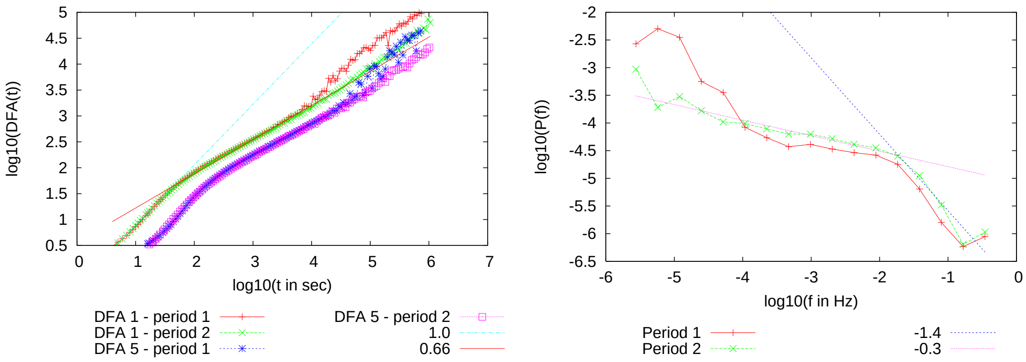

Fig. 5. Looking for long-range dependence using the DFA (left) and power spectrum method (right).

As we can see from Fig. (5) the changes in graphs of the DFA and power spectrum actually coincide. The crossovers happen at the same points in the frequency and time domain. The relation Eq. (9) also holds, though slightly worse for the range of 10 to 100 seconds. The overall conclusion is that there is very little evidence for long-range correlation in the data. Normal network traffic data shows a gradient of -1.0 in a frequency range of at least $10^{-0.5}$ to $10^{-3.0}$ Hz [10]. This is interesting as there is evidence that any network traffic that has shared a bottleneck router with self-similar traffic picks up the signature of that traffic [19]. This may be caused by the fact that we only observe the attempts of TCP connection but not the actual TCP traffic flow. It would be interesting to re-visit older traffic measurements and focus on the TCP SYN/ACK packets in the analysis.

4 The Sasser Worm Attack

In this section, we focus on two IP addresses that contact the telescope on port 445. This a port used by, for instance, the *Sasser* worm. One hallmark of the worm is that it tries to contact hosts repeatedly (twice) in a short time period at the same port. This is thought to be used to determine the operating system of the targeted host; it then follows up with specifically crafted exploit packets. Also, the Sasser worm tends to pick one class C network at a time and scan all hosts on this network. To estimate the time between attacks we compute the inter-arrival time for packets for each IP address inside the network telescope separately and neglect those smaller than 4 seconds. We plot the histogram of all these inter-arrival times, so as to estimate the PDF of the inter-arrival times. In Fig. (6), we see the result for two attackers in both observation periods. The total rate of attacks goes down from period 1 to 2. On either end of the PDF statistical errors are likely due to insufficient data. To get an idea of how a single attacker behaves we focus on one specific

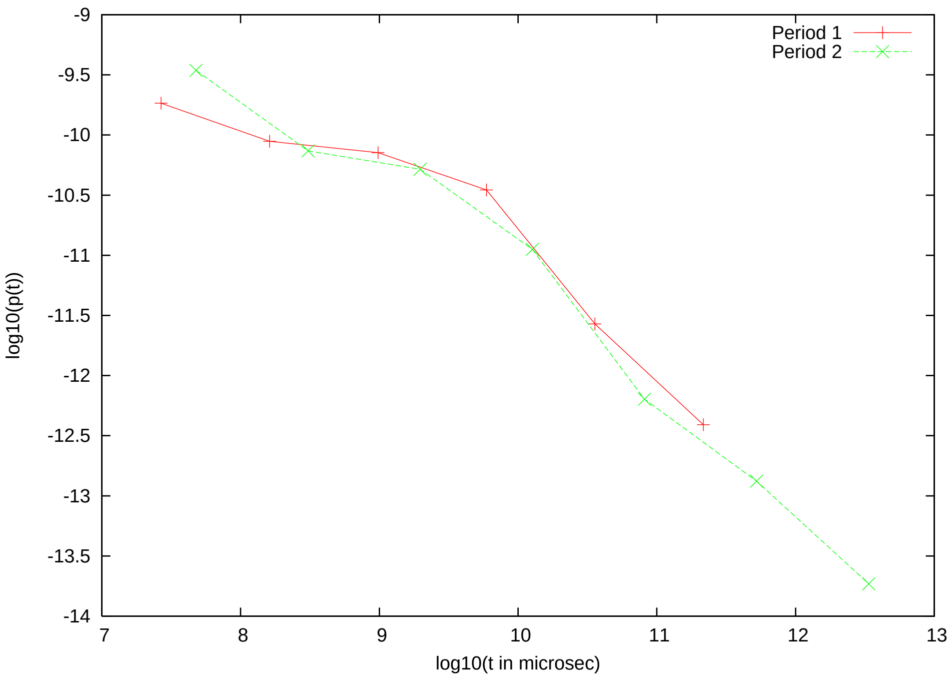

Fig. 6. inter-arrival time distribution for entire hostscans of the telescope by one particular attacker.

"attacker IP", and one IP number inside our network telescope. First we plot the aggregate number of packets sent out by that attacker, Fig. (7). Note that the attack rate has gone down in the second period from attacks every 3.8

hours to one in 24 hours. The inter-arrival times are clearly not exponential but are not heavy tailed either. The power spectrum shows a peak around the 24 hour mark, which indicates that the attacks arrive periodically. Similarly, the DFA plot hints a crossover at that point, which can be associated with periodic behaviour [16]. The values for the DFA and power spectrum match

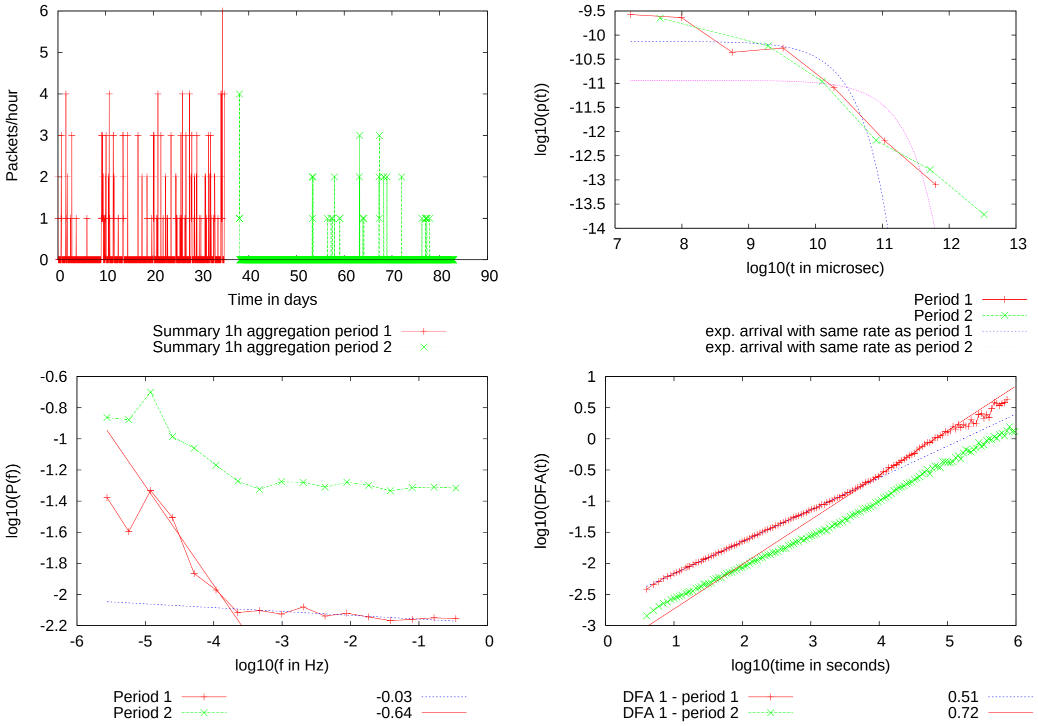

Fig. 7. Details of the behaviour of one particular attacking IP number. In the top row the packet rate plot and the inter-arrival time distribution of all packets from this attacker. In the bottom row the power spectrum and the DFA plot.

up fairly well according to Eq. (9). Again, there is no evidence for long-range dependence in the traffic trace. The gradient between 10^{-4} and 10^{-5} Hz may be more related to the peak indicating periodic events.

5 A denial-of-service attack

The telescope observes a significant number of denial-of-service attacks. An example is shown in Fig. (8). We can see how the packet rate rises sharply for a few hours. In addition, the inter-arrival times change dramatically, with a reasonable power law over several decades. Similarly the power spectrum and the DFA plot show a change. However, the change is simply a change from a signal with white noise to Brownian motion, The gradient of the power spectrum changes from $\delta = -1.8$ to $\delta = 0$, at the same time the DFA changes from $\alpha = 1.5$ to $\alpha = 0.5$, in line with Eq. (9). The crossover occurs in both cases at a timescale of about 500 seconds or 8.4 minutes. Though the packet summary data does actually have a visible trend, the different levels of the

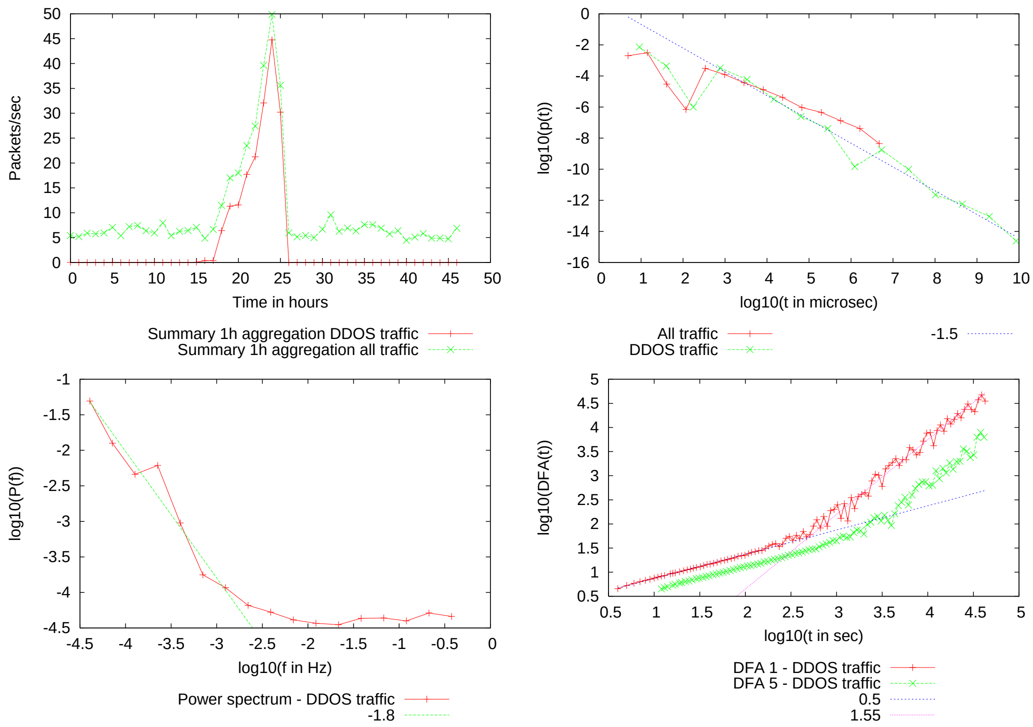

Fig. 8. Details of the backscatter from a denial-of-service attack. First the packet rate summary and the inter-arrival times. In the bottom row the power spectrum and DFA.

DFA show that these trends, at least for trends up to polynomials of degree 5, do not taint the correlation of the data. If the change from white noise to Brownian motion is not just caused by the trend in the data but by the additional attack itself we could use the inter-arrival time histogram, power spectrum and DFA plot to distinguish peaks caused by a simple rate increase and those caused by a DoS attack.

6 Sasser Worm: infection rules

In this simple model, we try to recreate the arrival of attacks and their associated packet traffic. To do this, we need to estimate a realistic scan rate for the attackers. The Sasser worm [8] uses 3 simple rules to choose a class C network for the next attack. If the infected host has got the IP number A.B.C.D

(i) a random number with probability 1/8

(ii) A.x.y.0 randomly with probability 1/2

(iii) A.B.x.0 randomly with probability 3/8

Interestingly we observe that the ratio of worm-infected to non-infected machines in either of those three regions above are of the same order of magnitude. We also know that the worm is likely to start 128 threads simultaneously for the attacks analysed in Section (4). So, we could make the following quick calculation:

Assume the attack rate, i.e. the rate class C networks are chosen at, is λ_A and the observed (by the Network telescope) attack rate is λ_O. The number of all possible class C networks is N_C, then the observed rate used by one attacker for a class C network telescope ought to be:

$$\lambda_A = N_C \lambda_O \tag{10}$$

assuming no preference for a particular subset. From our observations this model would predict rather high attack rates. For one particular attacker IP address we observe a rate of 5 attacks per day in the first and 1 attack per day in the second period. Given that there are 254^3 class C networks, this amounts to $\lambda_A = 1200$ attacks/sec or $\lambda_A = 300$ attacks/sec. According to [8], Sasser variants scan 510 to 40,960 IP addresses per second, meaning anything between 2 to 160 attacks per second. Notably, the observation in the first period is much higher than anything suggested in [8]. At closer inspection we find that the attacker IP address appears to be part of an address pool that is assigned via DHCP to ADSL subscribers. This might explain the higher scan frequency, as it corresponds to machine reboots or ADSL session lengths rather than the actual scan rates. In the second period we might simply have users that are connected for a longer time. The total number of observation of attacks in the first period is 223 and in the second is 45.

7 Conclusions

In this paper, we have investigated the nature of traffic seen by a passive network telescope. In our case the address of the network telescope (a class C network) has been unused for four years. Therefore virtually all of observed inbound traffic, amounting to 6 packets per second, are caused by malware such as worms and viruses. While hostscans (scans through many IP addresses in the network telescope for one specific port) are prevalent, portscans (scans of many ports on one IP address within the network telescope) are virtually non-existent.

The top 10 destination ports and top 100 IP-addresses account for virtually all the traffic, though both distributions show a very strong power law in the rank-frequency plot.

The vast majority of the traffic is TCP traffic, and we see no IPv6 traffic at all. The average packet rate shows no sign of long-term increase or decrease over several months. The traffic also shows surprisingly little sign of long-term correlation as is usually detected in network traffic. This could be caused by the lack of actual TCP traffic flows. To the best of our knowledge, this is the first time that anyone has specifically looked for long-range depencence in network traffic generated by malware. We also observe that the ratio of infected machines appears to be constant over A, B and C class networks.

With the onset of a denial-of-service attack, we saw that the signature of the traffic changed from white noise to Brownian motion, hinting at the possibility of distinguishing between background traffic noise increases and the introduction of possibly malicious traffic streams. With further investigation of this and future data, we hope to be able to produce realistic artificial packet workloads to test the performance of firewall designs. We will also investigate whether we can use the data in any way to map parts of the Internet, i.e. deduce its topology. In addition, we will investigate whether the observed attackers form any interesting network structures.

With the help of other subnets, we hope to gather data to form a distributed network telescope. This will be useful to see how attacks are correlated in time and (address) space.

Acknowledgments and Further Information

Uli Harder is funded by EPSRC (research grant PASTRAMI, GR/S24961/01). Jeremy Bradley is supported in part by the Nuffield Foundation under grant NAL/00805/G. The authors would like to thank Ashok Argent-Katwala for helpful suggestions for analysing and understanding the data.

The data used in this paper will be made public in due course on the AESOP website (http://aesop.doc.ic.ac.uk/); until then please email Uli Harder to arrange access to the data.

References

[1] D. Nicol, M. Liljenstam, and J. Liu, "Multiscale modeling and simulation of worm effects on the internet routing infrastructure," in *TOOLS'03, Proceedings of Computer Performance Evaluation: Modelling Techniques and Tools*, vol. 2794 of *Lecture Notes in Computer Science*, (University of Illinois at Urbana-Champaign), pp. 1–10, Springer-Verlag, September 2003.

[2] D. Moore, C. Shannon, G. M. Voelker, and S. Savage, "Network telescopes," tech. rep., CAIDA, 2003.

[3] F. Pouget, M. Dacier, and V. H. Pham, "Understanding threats: a prerequisite to enhance survivability of Computing Systems," in *IISW'04, International Infrastructure Survivability Workshop 2004, in conjunction with the 25th IEEE International Real-Time Systems Symposium (RTSS04) December 5-8, 2004 Lisbonne, Portugal*, December 2004.

[4] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson, "Characteristics of internet background radiation."

[5] V. Yegneswaran, P. Barford, and D. Plonka, "On the design and use of internet sinks for network abuse monitoring."

[6] W. E. Leland, M. Taqqu, W. Willinger, and D. Wilson, “On the self-similar nature of ethernet traffic (extended version),” *IEEE/ACM Transactions on Networking*, vol. 2, pp. 1–15, February 1994.

[7] H. Kim, I. Kang, and S. Bahk, “Real-time visualization of network attacks on high-speed links,” *IEEE Network*, pp. 30–39, September/October 2004.

[8] TrendMicro, “The Sasser Event: History and Implications,” tech. rep., www.trendmicro.com, June 2004.

[9] G. K. Zipf, *Human Behavior and the Principle of Least Effort.* Addison-Wesley, 1949.

[10] A. J. Field, U. Harder, and P. G. Harrison, “Measurement and modelling of self-similar traffic in computer networks,” *IEE Proceedings - Communications*, vol. 151, no. 4, pp. 355–363, 2004.

[11] U. Harder and M. Paczuski, “Correlated dynamics in human printing behavior,” tech. rep., Dept. of Computing, Imperial College London, December 2004. http://arxiv.org/abs/cs.PF/0412027.

[12] W. Press *et al.*, *Numerical Recipes in C 2nd ed.* CUP, 1993.

[13] A. L. Goldberger, L. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, “PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals,” *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000 (June 13). Circulation Electronic Pages: http://circ.ahajournals.org/cgi/content/full/101/23/e215.

[14] C. K. Peng, S. Havlin, H. E. Stanley, and A. L. Goldberger, “Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat time series,” *Chaos*, vol. 5, no. 1, pp. 82–87, 1995.

[15] C. K. Peng, S. V. Buldyrev, S. Havlin, M. Simons, H. E. Stanley, and A. L. Goldberger, “Mosaic organization of dna nucleotides,” *Phys. Rev. E*, vol. 49, pp. 1685–1689, 1994.

[16] J. W. Kantelhardt, E. Koscielny-Bunde, H. H. A. Rego, S. Havlin, and A. Bunde, “Detecting long-range correlations with detrended fluctuation analysis,” *Physica A*, vol. 295, pp. 441–454, 2001.

[17] J. Honerkamp, *Statistical Physics.* Springer, 2002.

[18] B. Pilgram and D. T. Kaplan, “A comparison of estimators of $1/f$ noise,” *Physica*, vol. D 114, pp. 108–122, 1998.

[19] A. Veres, K. S. Molnár, and G. Vattay, “On the propagation of long-range dependence in the internet,” *SIGCOMM Comput. Commun. Rev.*, vol. 30, no. 4, pp. 243–254, 2000.