A Baseline Numeric Analysis of Network Telescope Data for Network Incident Discovery

Bradley Cowie[1] and Barry Irwin[2]
Security and Networks Research Group
Department of Computer Science
Rhodes University
Grahamstown,South Africa
E-mail: [1]g06c5476@campus.ru.ac.za [2]b.irwin@ru.ac.za

Abstract— **This paper investigates the value of Network Telescope data as a mechanism for network incident discovery by considering data summarization, simple heuristic identification and deviations from previously observed traffic distributions. The datasets used for this analysis were obtained from a Network Telescope for the time period August 2005 to September 2009 which had been allocated a Class-C network address block at Rhodes University. The nature of the datasets were considered in terms of simple statistical measures obtained through data summarization which greatly reduced the processing and observation required to determine whether an incident had occurred. However, this raised issues relating to the time interval used for identification of an incident. A brief discussion into statistical summaries of Network Telescope data as "good" security metrics is provided. The summaries derived were then used to seek for signs of anomalous network activity. Anomalous activity detected was then reconciled by considering incidents that had occurred in the same or similar time interval. Incidents identified included Conficker, Win32.RinBot, DDoS and Norton Netware vulnerabilities. Detection techniques included identification of rapid growth in packet count, packet size deviations, changes in the composition of the traffic expressed as a ratio of its constituents and changes in the modality of the data. Discussion into the appropriateness of this sort of manual analysis is provided and suggestions towards an automated solution are discussed.**

Index Terms—**Network Telescope, incident discovery, numerical analysis**

I. Introduction

NETWORK Telescopes that utilize passive data collection [7] techniques provide researchers with a better understanding of anomalous activity occurring on the Internet through non-obtrusive analysis. This is due to the fact that all traffic received by a Network Telescope is unsolicited [12]. From this it can be inferred that all traffic collected must either be malicious, probing or caused by misconfiguration [12]. It should be noted that the class of malicious traffic is broad and includes: attempted viral or worm infection [13], DDoS (Distributed Denial of Service) [9], Network Spoofing [32] and Vulnerability Scanning. Due to the size of this class and the difficulties involved with the uncertainties in the classification process this paper shall only consider positive incident identification. It is plausible to consider that if a change in the composition of the network traffic has been identified then one of the aforementioned outcomes has occurred. The raw data consists of a large capture of packets. It is noted that it is difficult to convert this raw data into a usable measure. This is due to the issue of resolution, that is if too large a time interval is considered it becomes difficult to discern incidents due to interaction and dilution caused by other incidents. However if too small an interval is considered, far more processing is required to observe possible incidents. The danger of information overload then is possible due to the insignificance of the data and the bottleneck of having to consider each reading individually [33]. These issues are becoming particularly pertinent within the South African context due to the rapid increase of available bandwidth in South Africa. This is a result of the opening of new Internet routes to South Africa such as SEACOM [2] with future ventures including EASSy [1] and Main One [11]. Considering South Africa was rated as having the 7th highest number of Cyber Crimes perpetrators in 2009 [16], action needs to be taken lest South Africa become a hotbed of cybercriminal activity. This paper considers the idea of 'Network Incidents', these are large-scale points of interest where there is a considerable growth in malicious traffic directed towards the Internet in general or specific sections of the Internet. These 'incidents' usually require action such as security alerts or analysis from an official body such as CERT [8] or SANS [24]. The rampant spread of Conficker [30] is a noted example of a serious incident. From a telecommunications perspective being able to identify possible incidents as they occur is vital in order to react in an appropriate manner This is in order to avoid possible losses in revenue due to inflated network load. Furthermore, techniques discussed here could be extended to identify potential changes in user preferences in terms of traffic generated that is to be routed by a network provider.

A. Paper Organization

The remainder of the paper is structured as follows. Section II considers related work in the field of Network Telescope analysis. The data collection methodology, data summarization and summarized data as a Security Metric is discussed in Section III. Section IV discusses measuring normality within the context of Network Telescopes. A number of examples of

incident identification through manual analysis. are presented in Section V. Section VI discusses future work and the paper is concluded in section VII

II. Related works

A considerable amount of work has been conducted by the Information Security research community at large with regards to work in the field of Network Telescope analysis. In particular the researchers at CAIDA [3] (the Cooperative Association for Internet Data Analysis) have produced work defining the fundamentals of Network Telescopes [23]. Other work conducted by CAIDA includes observing large network incidents as they occur in particular Code Red Worm [21], SQLSlammer [20] and Witty Worm [26]. They have also developed frameworks for creating Distributed Network Telescope Nodes for the monitoring and analysis of network traffic on a global scale [10]. Work within Rhodes University has considered the relation between logical distances and packets collected by Network Telescopes [5], the graphical representations of network incidents through the use of InetViz a tool developed by van Riel and Irwin [29], [31] and mapping the Internet through space filling curves such as the Hilbert Curve [17]. While there has been worked conducted considering the statistical analysis of network traffic by clustering [19] and other means. Little work considers analysis by "simple" security metrics. This analysis acts as a stepping stone, allowing for future work into automated techniques for incident discovery. This is done by building upon the heuristics and observations obtained through manual analysis.

III. Data Summarization and Summarization as a Security Metric

Data summarization is an important process in the analysis of network traffic as it reduces the data set into more manageable components from which more meaningful analysis can be made. The original data set consists of over 33 million packets captured at Rhodes University during the time period August 2005 and September 2009 [6], [5]. This data was processed and imported into a SQL database consisting of entries containing the relevant components of each packet header. The packet payload was not included due to space and processing constraints. This data was then reduced to a smaller subset of numeric measures which provided a description of the data considering averages, medians, deviations and extrema. These statistics included Packet Counts, Summed Packet Size, Average Packet Size, Standard Deviations in Packet Size, Average TTL, Standard Deviation in TTL and Count per Hosts /32, /16 and /8. This data was grouped according to date at the hourly, daily, quarterly and yearly interval. The data was further subdivided according to type TCP, UDP and ICMP with subgroups by port for UDP and TCP and by ICMP type for ICMP. These summarizations are now considered as security metrics. According to the *Guide to Security* [23], a seminal paper in the field of security metrics, "Good metrics are those that are SMART, i.e. specific, measurable, attainable, repeatable, and time-dependent". The analysis provided from the aforementioned measures meet these criteria of good security metrics as measures of security incidents. This is due to the fact they are specific to a network incident such as IP Spoofing or viral activity. Furthermore, these statistics provide a measurable analysis as it can be inferred whether an incident has occurred by considering deviations from normality. In conclusion the results are clearly attainable and repeatable and the measures vary dependent on the time and time resolution observed. It thus follows that these summarizations are good security metrics according to [23].

IV. Describing Normal Traffic for a Network Telescope

In order to detect deviations from normal traffic a measure of normal traffic is required. This is a difficult concept to succinctly define within the realm of network traffic analysis. This is caused by the vast quantity of traffic that is to be analyzed and classified, dilution of incidents due to interplay with other incidents and frequent changes in both the physical and logical composition of the Internet. However it should be noted, as was previously alluded to, the traffic observed by a Network Telescope is fundamentally different to the traffic seen by other nodes connected to the Internet. Firstly, all traffic obtained is unsolicited and secondly because the telescope used is passive, the TCP handshake cannot by completed, it thus follows that all packets with SYN and SYN-ACK flags set are done so accidentally or through malicious intent [15]. Normality of network traffic can be decomposed into a number of measures of the normality of network activity. Packet Count per port, Ratio of Packet per port, Packet Size distribution and Packet Proportion per Host Network are a few examples of such norms. The problem with some of these measure, such as mean and standard deviation, is that they are easily influenced by outliers in the data set. To solve this more robust indicators such as median and expected deviation are used, however it is noted that they are more computationally expensive to calculate. Table I shows selected means and medians that describe the normality of the data for monthly intervals for port 5678. From this table it can be seen that mean values in the case of Packet Count and Packet Size have been influenced by some large values whereas the median gives a more better estimate of normality.

In the case of major incidents, which result in a significant change in traffic distribution, the normality of the data is skewed. This makes it difficult measure future incidents as they would be measured relative to the skewed normality. A solution to this would be to remove the anomalous activity from the data set. For example, as shown in Figure 1, it is clear pre-Conficker, that is before 21 October 2008, there is a clear relation between the packet counts of traffic inclusive and exclusive of port 445. That is the traffic peaks and dips at relatively the same place and the general shapes are similar. However afterwards there was a rapid growth in traffic inclusive of port 445. In fact the traffic inclusive of port 445 shows exponential growth while the traffic exclusive of port 445 continues in a similar fashion as previously observed.

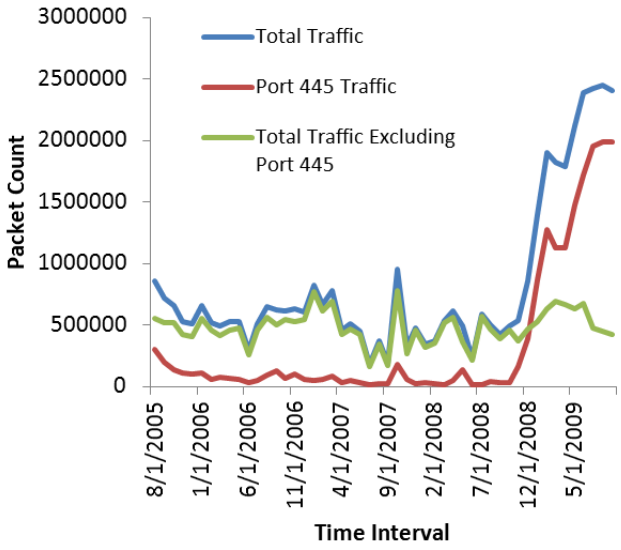

Figure 1. Plot of Packet Counts per month between August 2005 to September 2009, Highlighting the significant effect Conficker had on the distribution of anomalous activity collected from the telescope.

Quantity	Median	Mean
Packet Count	0	2215.5
TTL	91	108
Packet Weight	5	583893

Table I
TABLE OF MEANS AND MEDIANS FOR PORT 5678 GROUPED AT A MONTHLY LEVEL FOR 2005

V. INCIDENT DISCOVERY

This section considers some simple rules or heuristics derived from observation and theory that could be used to detect anomalous traffic.

A. Analysis by ratio of an observed quantity

A ratio of the total observed count and a particular quantity provides a useful and scalable measure of the traffic distribution. Whereas a plain count is not sensitive to changes that occur should there be a sudden increase in traffic flow, a ratio adapts to these changes adequately. In particular, a ratio of top quantities are of interest as these show a clear variation in traffic composition

1) Variation in Packet Count Ratios: It is expected that "top ports" will constitute the majority of packet counts. However some degree of normality is expected within these top ports assumed from the Central Limit Theorem [14]. Thus it can be concluded that major shifts in these ratios are indicative of anomalous behavior. Mathematically this rule could be described as $P/n > c \times P_{average}/n$ where n is the sample size considered at a time interval, c an empirically derived constant which is context sensitive, P is the observed count and $P_{average}$ is the average count for previously observed intervals. From Figure 2 it is observed that in 2009 approximately 73% of all traffic received was destined to port 445. Clearly this exceeds the previous average for port 445 and considering the time period this anomaly is clearly caused by

Date	Packet Count	Percentage of Total Count
2005	2	0.0037
2006	1	0.0018
2007	0	0
2008	4390	99.9945

Table II
TABLE OF PACKET COUNTS AND PERCENTAGE OF TOTAL COUNT FOR PORT 5678 GROUPED AT A YEARLY LEVEL

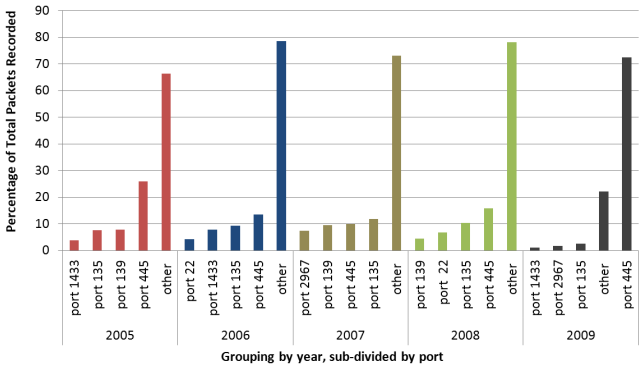

Figure 2. Percentage of total traffic for a given port grouped on a yearly scale. This graph shows a major shift in the distribution bewteen 2008-2009, that is the majority of traffic recorded is bound for port 445.

Conficker. Again the issue of time resolution is present as a yearly count identifies the incident in 2009 and not late 2008.

2) Distribution of top port counts: It would be expected that in a normal traffic distribution that the top ports would have packet counts that are well distributed over the entire period. If this were not the case it is possible that some anomalous activity has taken place. Table II provides measures of port 5678, which had been identified as one of the top 40 ports in the dataset. It is clear from this table that most of this data lies within one specific year and may be indicative of anomalous behavior. Researching relevant trusted security websites yields exploits that existed in Symantec Netware on port 5678 [27], [4] during early 2008.

B. Analysis by Deviation in Packet Size and ICMP Type

The standard deviation, σ, of a measured quantity provides a sense of the spread of the measured quantity. While it is computationally easy to calculate, relatively speaking, it is not a robust measure, that is, it is subject to be influenced by large outliers. While this may be inappropriate when dealing with identifying norms, large outliers are exactly what may tip an analyst off to possible incidents. As previously mentioned, from the Central Limit Theorem, it is expected that approximately 95% of all data to lie within two standard deviations of the mean. The above considerations allow us to construct a useful rule in terms of variation. That is, if there is a sudden change in the deviation of a quantity for a sufficiently large number of samples, then an incident has most likely occurred. Mathematically this can be expressed as $\sigma > 2 \times c \times \sigma_{old}$ where σ is the measured standard deviation, c is an empirically derived constant and σ_{old} is the previous measurement for the standard deviation. This measure is useful

Date	Average Packet Size
February 19, 2007	61.57
February 20, 2007	69.12
February 21, 2007	60.88
February 22, 2007	60.93

Table III
TABLE OF PACKET SIZE AVERAGES GROUPED AT A DAILY LEVEL

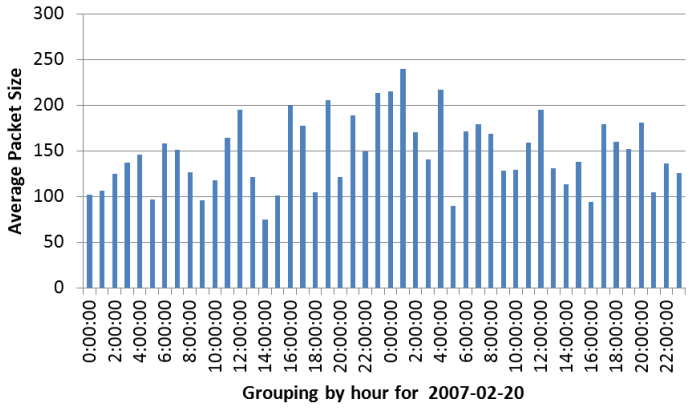

Figure 3. Plot depicting the average packet size for port 2967 for 20 February 2007 grouped at an hourly level. It is clear to see there is a large amount of variation in packet size which was atypical of previous readings.

Month	Packet Count	Std. Dev of Packet Size
December 2006	159675	0.66
January 2007	144334	2.93
February 2007	104435	4.19
March 2007	96008	2.78
April 2007	20062	2.83
May 2007	16666	1.96
June 2007	21828	2.98
July 2007	25812	1.69
August 2007	7323	4.07

Table IV
TABLE OF PACKET COUNT AND STANDARD DEVIATION OF PACKET SIZE MEASURED AT MONTHLY INTERVALS

Day	Packet Count	Std. Dev of Packet Size
February 16, 2007	3142	0.23
February 17, 2007	14029	0.85
February 18, 2007	3423	0.88
February 19, 2007	5802	1.48
February 20, 2007	21563	57.37
February 21, 2007	3006	1.88
February 22, 2007	3507	1.61
February 23, 2007	5049	1.00

Table V
TABLE OF PACKET COUNT AND STANDARD DEVIATION OF PACKET SIZE MEASURED AT DAILY INTERVALS

in incident discovery if it is considered that some viruses randomly pad the payload of their packets in an attempt to prevent their packets from being filtered out by a firewall and making it more difficult to design IDS (Intrusion Detection System) signatures to match the viral traffic. A sharp change in the amount of traffic on a specific port or ICMP type coupled with a change in standard deviation in packet size is a good numeric description of this sort of anomalous behavior

1) Variation in Packet Size: For this section port 2967 will be considered as a port of interest. The minimum size of a packet assuming an Ethernet frame is 64 bytes. In general it is expected that most packets to just above this minimum size as illustrated in Table III. However averages in general make poor measures of variability unless the time resolution is sufficiently fine-grained. Considering the ports with the highest variation in 2007 a number of potential incidents are considered. If the deviations at monthly level are considered for port 2967 there does not appear to be any significant difference in the standard deviation as shown in Table IV. However, considering a daily standard deviation reveals 20 February 2007 to be a day of interest as shown in Table V. Plotting the deviation at an hourly level shows the considerable deviation in packet sizes during this time period as shown in Figure 3. This illustrates the point that the time interval considered is highly significant. An investigation from reputable Internet security sources yields that W32.Rinbot [28] started to emerge in the wild as early as February 2007. It exploited vulnerabilities in applications running on port 2967 through specially crafted RPCs (Remote Procedure Calls) with randomized packet size padding.

C. Analysis by counts

Counts of data grouped according to some critera provides a simple and fast measure of normality by comparison with previously measured means. It should be noted that this sort of measure is prone to be ineffective should there be a significant change in the logical or physical network topology of the Network Telescope.

1) Cases of denial of service through spoofing: A common technique used in DDoS is to spoof an IP range and then use this range to attack a server. This of course, assuming sufficient load, causes the server to incidentally stop responding resulting in times outs. These time outs result in the generation of ICMP Type 11 packets which are sent back to the spoofed address space. Occasionally it occurs that this address space actually belongs to a Network Telescope. Table VI alludes to the possibility of this sort of activity occurring February of 2009. After considering tables representing packet counts for ICMP type 11 at a daily resolution it was determined that some form of DDoS could potentially have been perpetrated during 17 - 18 February 2009 as shown in Figure 4.

Date	Packet Count
November 2008	1195
December 2008	1722
January 2009	2962
February 2009	163213
March 2009	2142

Table VI
TABLE OF ICMP TYPE 11 PACKET COUNT BY MONTH

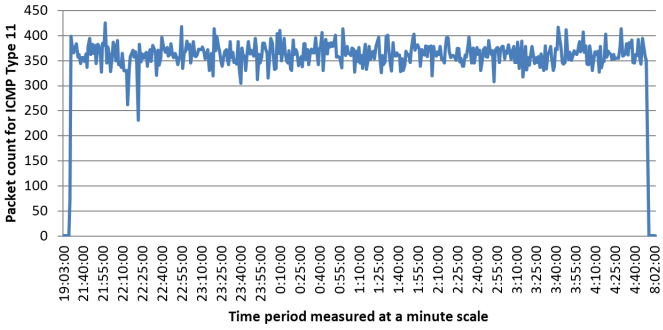

Figure 4. Plot depicting the sudden apperance of ICMP Type 11 Traffic during the 17th and 18th of February 2008.

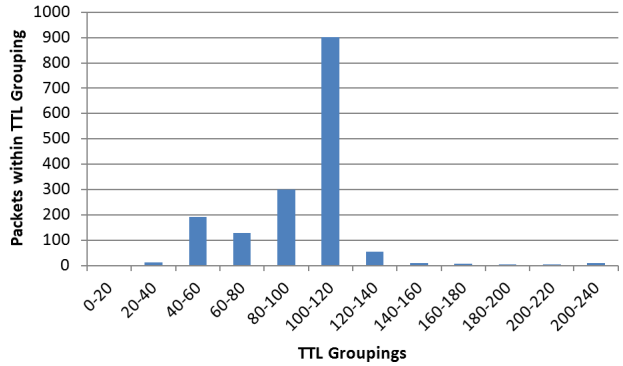

Figure 5. Bins of Average TTL values measured at the time interval of a month. There appears to be 4 possibly 3 distinct modes to the data.

D. Analysis by modality

Some quantities have clear grouping of data around certain values known as modes. Variation in the existence of these modes and the points where values cluster around TTL (Time To Live) provides an interesting modal behavior as shown in Figure 5. Using these modalities it may be possible to detect variations in logical topology, potential spoofing and supplement OS identification techniques from network packets.

E. Issues with Manual Analysis

The process of identifying events and devising heuristics is a time consuming process and lacks completeness in the analysis. To carefully consider 65536 ports is no small feat and this is compounded by the fact that an analysis of network traffic that relies on ports alone is very limited in what can be concluded. As has been identified in this section, the time interval at which events are considered greatly affects the outcome of the analysis, considering very fine time scales is particularly costly in terms of processing and time required for analysis. Further, this sort of analysis relies on the fact that incidents are well documented and this information is available, which is often not the case.

VI. Future Work

The authors indent to extend this research in baseline analysis into more automated techniques to be used in conjunction with manual techniques. This will be done by considering Artificial Intelligence (AI) techniques. These proposed techniques for automation include investigations into :

- Constructing incident listings in both an automated and manual sense. One potential automated technique is to build a web-crawler that collects data from reputable sources such the US CERT, The Internet Storm Center and SANS. The incident shall then be classified according to the anomalous activity that defines it. For example, Conficker could be defined as rapid growth in port 445 affecting Windows machines.
- Optimization of heuristic rules through Genetic Algorithms for effective incident discovery. This sort of research would consider ways to artificially breed rules for incident discovery. This is in order to maximize the number of incidents detected while simultaneously minimizing the number of false-negatives. This sort of work is reminiscent of the work done by Nottingham and Irwin in breeding rules for fast packet classification [22].
- Aggregating data collected by multiple Network Telescope nodes to allow for more distributed monitoring.
- The use of marked incident-time series data to train Neural classifiers to detect anomalous activity. This would include an investigation into which AI construct is most appropriate of Fuzzy Logic, Neural Networks and Bayesian Networks.

At a lower level, further investigation into useful measures is also required. These include :

- Sliding averages, these are values that consider the data for a certain period of time and change appropriately according to the context.
- Considering more advanced statistical techniques such as ANOVA and regression analysis as good security metrics for incident identification.
- Using incident modeling to match incidents to a particular modeled incident type. Some basic work could involve considering the SIRS model [18] and other models of virus [25] propagation. This is in order to generate simulated network traffic to be used as training information for neural classifiers.

VII. Conclusion

This paper has considered a number of measures, derived by data summarization, which have been shown to be competent in the area of network incident detection. The way in which these values were obtained and the appropriateness of these measures as Security Metrics has been discussed. The main measures considered included quantity as ratios of quantities, simple counts, standard deviations in quantities and analysis by modality. These techniques detected the following incidents: Conficker, W32.Rinbot, DDoS and attempted exploits in Norton Netware. The time interval at which analysis took place was identified as an important factor in this sort of analysis with examples of this provided. The weaknesses of manual analysis was described and techniques for more automated identification were discussed.

ACKNOWLEDGEMENTS

The work reported in this paper was undertaken in the Telkom Centre of Excellence in Distributed Multimedia at Rhodes University, with financial support from Telkom, Comverse, Tellabs, Stortech, Amatole Telecom Services, Bright Ideas 39 and THRIP.

REFERENCES

[1] Eastern Africa Submarine Cable System (EASSy). [Online]. Available: http://www.eassy.org/. [Accessed: Apr 20, 2010].

[2] SEACOM - South Africa - East Africa - South Asia - Fiber Optic Cable. [Online]. Available: http://www.seacom.mu/. [Last Accessed: 18 April, 2010].

[3] The Cooperative Association for Internet Data Analysis. [Online]. Available: http://www.caida.org/home/. [Accessed: 21 April,2009].

[4] Vulnerability Summary for CVE-2008-1701. [Online]. Available : http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2008-1701, Accessed.

[5] R. Barnett and B. Irwin. An analysis of logical network distance on observed packet counts for network telescope data. In *SATNAC 2009*, 2009.

[6] Nick Pilkington Barry Irwin and Blake Friedman. A geopolitical analysis of long term internet network telescope traffic. In *SATNAC*, 2007.

[7] CAIDA. Passive Data Collection: UCSD Network Telescope. [Online]. Available: http://www.caida.org/data/passive/network_telescope.xml. [Accessed: April 19, 2010], January 2010.

[8] CERT. Carnegie Mellon University's Computer Emergency Response Team. [Online]. Available: http://www.cert.org/. [Accessed: Apr 20, 2010].

[9] Kun chan Lan, Alefiya Hussain, and Debojyoti Dutta. Effect of Malicious Traffic on the Network, 2003.

[10] Kimberly Claffy, Young Hyun, Ken Keys, Marina Fomenkov, and Dmitri Krioukov. Internet Mapping: From Art to Science. In *CATCH '09: Proceedings of the 2009 Cybersecurity Applications & Technology Conference for Homeland Security*, pages 205–211, Washington, DC, USA, 2009. IEEE Computer Society.

[11] Main One Cable Company. Main One Cable. 2009. [Online]. Available: http://www.mainonecable.com/, [Accessed: Apr. 20, 2010].

[12] G. Voelker St. Savage D. Moore, C. Shannon. Network Telescopes: Technical Report. [Online]. Available: http://www.caida.org/publications/papers/2004/tr-2004-04/tr-2004-04.pdf. [Accessed: 10 April, 2010], 2004.

[13] OECD Organisation for Economic Co-operation and Development. *Computer Viruses and Other Malicious Software: A Threat to the Internet Economy*. OECD, 2009.

[14] C. Grinstead and J. Snell. *Grinstead and Snell's Introduction to Probability*. American Mathematical Society, 4th edition, July 2006.

[15] Uli Harder, Matt W. Johnson, Jeremy T. Bradley, and William J. Knottenbelt. Observing Internet Worm and Virus Attacks with a Small Network Telescope. *Electron. Notes Theor. Comput. Sci.*, 151(3):47–59, 2006.

[16] Internet Crime Compliant Center (IC^3). 2009 Internet Crime Report. [Online]. Available: http://www.ic3.gov/media/annualreport/2009_IC3Report.pdf. [Accessed: Apr 21, 2010], 2010.

[17] Barry Irwin and Nick Pilkington. High Level Internet Scale Traffic Visualization Using Hilbert Curve Mapping. In *VizSEC*, pages 147–158, 2007.

[18] Qiming Liu, Rui Xu, and Shaojie Wang. Modelling and Analysis of an SIRS Model for Worm Propagation. In *CIS '09: Proceedings of the 2009 International Conference on Computational Intelligence and Security*, pages 361–365, Washington, DC, USA, 2009. IEEE Computer Society.

[19] David Marchette. A statistical method for profiling network traffic. In *Proceedings of the Workshop on Intrusion Detection and Network Monitoring*, pages 119–128, Berkeley, CA, USA, 1999. USENIX Association.

[20] David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, and Nicholas Weaver. Inside the Slammer Worm. *IEEE Security and Privacy*, 1(4):33–39, 2003.

[21] David Moore, Colleen Shannon, and k claffy. Code-Red: a case study on the spread and victims of an internet worm. In *IMW '02: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment*, pages 273–284, New York, NY, USA, 2002. ACM.

[22] Alastair Nottingham and Barry Irwin. GPU packet classification using OpenCL: a consideration of viable classification methods. In Barry Dwolatzky, Jason Cohen, and Scott Hazelhurst, editors, *SAICSIT Conf.*, ACM International Conference Proceeding Series, pages 160–169. ACM, 2009.

[23] Shirley Payne. A Guide to Security Metrics, SANS Institute. pages 1–2, 2002.

[24] SANS. SANS Internet Storm Center; Cooperative Network Security Community - Internet Security. [Online]. Available: http://isc.sans.org/. [Accessed : 20 April, 2010].

[25] Giuseppe Serazzi and Stefano Zanero. Computer Virus Propagation Models. In *In Tutorials of the 11th IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems*. Springer-Verlag, 2003.

[26] Colleen Shannon and David Moore. The Spread of the Witty Worm. *IEEE Security and Privacy*, 2(4):46–50, 2004.

[27] Symantec. MSRPC Client Svc Netware DoS. [Online]. Available: http://www.symantec.com/business/security_response/attacksignatures/detail.jsp?asid=21881. [Accessed : 17 April, 2010], 2007.

[28] Symantec. W32.rinbot.a. [Online]. Available http://www.symantec.com/security_response/writeup.jsp?docid=2007-021615-1555-99. [Accessed : 17 April,2010], February 2007.

[29] Jean-Pierre van Riel and Barry Irwin. InetVis, a visual tool for network telescope traffic analysis. In *Afrigraph*, pages 85–89, 2006.

[30] US-CERT. Conficker worm targets microsoft windows systems. [Online] : http://www.us-cert.gov/cas/techalerts/TA09-088A.html, April 2009.

[31] Jean-Pierre van Riel and Barry Irwin. Identifying and Investigating Intrusive Scanning Patterns by Visualizing Network Telescope Traffic in a 3-D Scatter-plot. In *ISSA*, pages 1–12, 2006.

[32] Victor Velasco. Introduction to IP Spoofing. SANS Institute InfoSec Reading Room:2–3, 2000.

[33] David D. Woods, Emily S. Patterson, and Emilie M. Roth. Can We Ever Escape from Data Overload? A Cognitive Systems Diagnosis. *Cognition, Technology & Work*, 4(1):22–36, 2002.

Bradley Cowie Bradley Cowie has an BSc(Hons) in Computer Science from Rhodes University. This BSc(Hons) was obtained whilst under the supervision of Barry Irwin. Bradley is currently working towards his MSc in Computer Science. Bradley's main research interests are network analysis and the application of Artifical Intelligence to trend monitoring.

Barry Irwin Mr Barry Irwin is currently completing his PhD research on the use o Passive Sensors as a means for inferring hostile network activity on the Internet. He holds a MSc in Computer Science from Rhodes University, and has research interests in Information Security to modern IP networks particularly: adaptive and automated network defence, early warning systems, and data visualisation techniques to support these.