

# AN EXPLORATORY STUDY OF TECHNIQUES IN PASSIVE NETWORK TELESCOPE DATA ANALYSIS

Submitted in fulfilment  
of the requirements for the degree of

MASTER OF SCIENCE

of Rhodes University

Bradley Cowie

*Grahamstown, South Africa*

November 2012

## **Abstract**

Careful examination of the composition and concentration of malicious traffic in transit on the channels of the Internet provides network administrators with a means of understanding and predicting damaging attacks directed towards their networks. This allows for action to be taken to mitigate the effect that these attacks have on the performance of their networks and the Internet as a whole by readying network defences and providing early warning to Internet users. One approach to malicious traffic monitoring that has garnered some success in recent times, as exhibited by the study of fast spreading Internet worms, involves analysing data obtained from network telescopes.

While some research has considered using measures derived from network telescope datasets to study large scale network incidents such as Code-Red, SQLSlammer and Conficker, there is very little documented discussion on the merits and weaknesses of approaches to analyzing network telescope data. This thesis is an introductory study in network telescope analysis and aims to consider the variables associated with the data received by network telescopes and how these variables may be analysed. The core research of this thesis considers both novel and previously explored analysis techniques from the fields of security metrics, baseline analysis, statistical analysis and technical analysis as applied to analysing network telescope datasets. These techniques were evaluated as approaches to recognize unusual behaviour by observing the ability of these techniques to identify notable incidents in network telescope datasets.

## Acknowledgements

A number of individuals and organizations must be thanked for their contribution to this thesis. Without the knowledge and resources provided by these contributors this thesis would have been an arduous task to complete.

Thanks is given to the authors parents for their support during the completion of this thesis. The constant encouragement, patience and positivity expressed helped ensure the success of this thesis.

Dr. Barry Irwin is thanked for directing this thesis from the stages of conceptualization, design, implementation and finally towards completion. Dr. Irwin's deep understanding of network monitoring and experience with network telescopes provided many of the concepts discussed in this document.

Thanks must be given to CAIDA and Team Cymru, industry experts in the field of network telescope data analysis, for their contribution and insight in the field of network telescope analysis. Special thanks is given to Stephen Gill for correspondence with regards to the approaches used by Team Cymru for identifying potential incidents in their datasets.

Recognition must be given for the feedback and interest taken by the staff and students of the computer science department at Rhodes University. Comments and critiques provided by members of the department were taken under advisement and have had a significant effect on the direction of the work produced.

This work was undertaken in the Distributed Multimedia CoE at Rhodes University, with financial support from Telkom SA, Stortech, Tellabs, Easttel, Bright Ideas 39 and THRIP. Thanks is given to Telkom SA for the funding provided to the author allowing for the completion of this thesis.

## ACM Computing Classification System Classification

Thesis classification under the ACM Computing Classification System (1998 version, valid through 2012).

### **Primary Classifications**

*C.2.0[General]* : Security and protection.

### **Additional Classifications**

*C.2.3[Network Operations]* : Network monitoring.

*C.2.3[Network Operations]* : Network management.

# Table of Contents

|   |             |
|---|-------------|
| <b>List of Figures</b>                                    | <b>vi</b>   |
| <b>List of Tables</b>                                     | <b>viii</b> |
| <b>1 Introduction</b>                                     | <b>1</b>    |
| 1.1 Problem Statement . . . . .                           | 2           |
| 1.2 Research Objectives . . . . .                         | 3           |
| 1.3 Research Approach . . . . .                           | 4           |
| 1.4 Research Scope . . . . .                              | 4           |
| 1.5 Document Conventions . . . . .                        | 5           |
| 1.6 Document Outline . . . . .                            | 6           |
| <b>2 Literature Review</b>                                | <b>8</b>    |
| 2.1 Networking Concepts and Internet Monitoring . . . . . | 9           |
| 2.1.1 Types of Network Traffic . . . . .                  | 9           |
| 2.1.2 Motivation for Network Security . . . . .           | 10          |
| 2.2 Network Telescope Concepts . . . . .                  | 12          |
| 2.2.1 Properties of Network Telescopes . . . . .          | 12          |
| 2.2.2 Limitations of Network Telescopes . . . . .         | 14          |
| 2.3 Case Studies . . . . .                                | 14          |
| 2.3.1 CAIDA . . . . .                                     | 15          |

---

|          |   |           |
|----------|---|-----------|
| 2.3.2    | Team Cymru . . . . .                              | 17        |
| 2.4      | Data Analysis Techniques and Approaches . . . . . | 17        |
| 2.4.1    | Security Measurements and Metrics . . . . .       | 18        |
| 2.4.2    | Baselines . . . . .                               | 18        |
| 2.4.3    | Moving Averages . . . . .                         | 19        |
| 2.4.4    | Technical Analysis . . . . .                      | 19        |
| 2.4.5    | Descriptive Statistics . . . . .                  | 23        |
| 2.5      | Summary . . . . .                                 | 23        |
| <b>3</b> | <b>Dataset Overview</b>                           | <b>25</b> |
| 3.1      | Dataset Origins . . . . .                         | 26        |
| 3.2      | Decomposing Network Traffic . . . . .             | 27        |
| 3.3      | Database Structure and Summarization . . . . .    | 30        |
| 3.4      | Descriptive Statistics . . . . .                  | 30        |
| 3.4.1    | Destination IP . . . . .                          | 30        |
| 3.4.2    | Source IP . . . . .                               | 32        |
| 3.4.3    | TTL . . . . .                                     | 33        |
| 3.4.4    | Packet Size . . . . .                             | 34        |
| 3.4.5    | Source Port . . . . .                             | 36        |
| 3.4.6    | Destination Port . . . . .                        | 37        |
| 3.4.7    | ICMP Type . . . . .                               | 39        |
| 3.4.8    | OS Type . . . . .                                 | 40        |
| 3.4.9    | Variable Frequency . . . . .                      | 41        |
| 3.5      | Summary . . . . .                                 | 41        |
| <b>4</b> | <b>Notable Incidents</b>                          | <b>43</b> |

---

|          |   |           |
|----------|---|-----------|
| 4.1      | DDoS Backscatter . . . . .                        | 44        |
| 4.2      | Conficker . . . . .                               | 46        |
| 4.3      | Windows Net Send Malware . . . . .                | 47        |
| 4.4      | Scanning Activity . . . . .                       | 48        |
| 4.5      | SQLSlammer . . . . .                              | 49        |
| 4.6      | Summary . . . . .                                 | 50        |
| <b>5</b> | <b>Metrics for Incident Identification</b>        | <b>52</b> |
| 5.1      | Normality for Network Telescope Metrics . . . . . | 53        |
| 5.2      | Analysis by Counts . . . . .                      | 54        |
| 5.2.1    | Total Packet Count . . . . .                      | 54        |
| 5.2.2    | Counts by Specific Packet Sizes . . . . .         | 55        |
| 5.2.3    | Destination IP Counts . . . . .                   | 56        |
| 5.2.4    | Port and Type Counts . . . . .                    | 57        |
| 5.3      | Analysis by Aggregation . . . . .                 | 58        |
| 5.3.1    | Average Packet Size . . . . .                     | 58        |
| 5.3.2    | Average TTL Values . . . . .                      | 59        |
| 5.3.3    | Average Counts per Source IP . . . . .            | 59        |
| 5.3.4    | Average Counts per Destination IP . . . . .       | 60        |
| 5.3.5    | Maximum Packets per Destination IP . . . . .      | 61        |
| 5.4      | Analysis by Groupings . . . . .                   | 61        |
| 5.4.1    | Protocol Counts . . . . .                         | 61        |
| 5.4.2    | OS Grouping . . . . .                             | 62        |
| 5.4.3    | Bogon Counts . . . . .                            | 63        |
| 5.5      | Summary . . . . .                                 | 64        |

---

|          |   |           |
|----------|---|-----------|
| <b>6</b> | <b>Incident Identification Techniques</b>   | <b>66</b> |
| 6.1      | Baselines Analysis . . . . .                | 67        |
| 6.1.1    | Specific-time Baselines . . . . .           | 67        |
| 6.1.2    | Recurring Baselines . . . . .               | 68        |
| 6.1.3    | Rolling Baselines . . . . .                 | 68        |
| 6.2      | Technical Analysis . . . . .                | 70        |
| 6.2.1    | Donchian Price Channels . . . . .           | 70        |
| 6.2.2    | Bollinger Bands . . . . .                   | 72        |
| 6.2.3    | Keltner Channels . . . . .                  | 73        |
| 6.2.4    | Moving Average Envelopes . . . . .          | 74        |
| 6.3      | Summary . . . . .                           | 75        |
| <b>7</b> | <b>Results</b>                              | <b>76</b> |
| 7.1      | Metrics . . . . .                           | 77        |
| 7.1.1    | Count-Based Metrics . . . . .               | 77        |
| 7.1.2    | Aggregated Metrics . . . . .                | 80        |
| 7.1.3    | Grouped Metrics . . . . .                   | 83        |
| 7.2      | Identification Aids . . . . .               | 85        |
| 7.2.1    | Baseline Analysis . . . . .                 | 86        |
| 7.2.2    | Technical Analysis . . . . .                | 88        |
| 7.3      | Recommendations . . . . .                   | 91        |
| 7.4      | Summary . . . . .                           | 92        |
| <b>8</b> | <b>Conclusion</b>                           | <b>93</b> |
| 8.1      | Summary of Previous Chapters . . . . .      | 93        |
| 8.2      | Review of the Research Objectives . . . . . | 94        |

---

|          |   |            |
|----------|---|------------|
| 8.3      | Limitations of Work . . . . .                   | 95         |
| 8.4      | Future Work . . . . .                           | 96         |
| <b>A</b> | <b>Glossary</b>                                 | <b>110</b> |
| <b>B</b> | <b>Networking Concepts</b>                      | <b>114</b> |
| B.1      | Bogon IP List . . . . .                         | 114        |
| B.2      | Common Ports and ICMP Types . . . . .           | 114        |
| B.3      | Operating Systems Identified from p0f . . . . . | 116        |

# List of Figures

|     |   |    |
|-----|---|----|
| 2.1 | Components of a Network Telescope . . . . .                 | 12 |
| 2.2 | CoralReef Monitoring System (Chicago B Monitor) . . . . .   | 16 |
| 2.3 | Financial Stockchart of SPY with TA Indicators. . . . .     | 21 |
| 3.1 | ERD of Significant Packet Trace Variables. . . . .          | 30 |
| 3.2 | Histogram of Destination IP Counts. . . . .                 | 31 |
| 3.3 | Histogram of TTL Counts. . . . .                            | 34 |
| 3.4 | Histogram of Packet Size. . . . .                           | 35 |
| 3.5 | Histogram of ICMP Type Counts. . . . .                      | 39 |
| 4.1 | Counts for ICMP type 11 during DDoS.D (/minute). . . . .    | 46 |
| 4.2 | Packet Counts for Dest. Port 445/tcp (/hour). . . . .       | 46 |
| 4.3 | A Hoax Net Send Message. . . . .                            | 48 |
| 4.4 | Packet Counts for Dest. Ports 1026-1027/udp (/day). . . . . | 48 |
| 4.5 | Packet Counts for Dest. Port 1434/udp (/day). . . . .       | 50 |
| 5.1 | Total Packet Counts (/month) . . . . .                      | 54 |
| 5.2 | Total Packet Counts (/day) . . . . .                        | 55 |
| 5.3 | Packets Counts per Packet Size (/day). . . . .              | 56 |
| 5.4 | Packets Counts per Destination IP (/day). . . . .           | 56 |
| 5.5 | Packets Counts for Dest. Port 445/tcp (/day). . . . .       | 57 |

---

|      |  |    |
|------|--|----|
| 5.6  | Top Five TCP Ports (/year).  | 58 |
| 5.7  | Average Packet Size (/day).  | 58 |
| 5.8  | Average TTL Value (/day).  | 59 |
| 5.9  | Average Packet Count per Source IP (/day).                                 | 60 |
| 5.10 | Average Packet Count per Destination IP (/day).                            | 60 |
| 5.11 | Maximum Packets per Destination IP (/day).                                 | 61 |
| 5.12 | Packets Counts per Protocol (/day).  | 62 |
| 5.13 | Packet Counts per OS Grouping (/day).                                      | 63 |
| 5.14 | Bogon IP Counts (/day).  | 64 |
|      |  |    |
| 6.1  | One-Year Recurring Baseline of Total Traffic (/day).                       | 68 |
| 6.2  | Six-Month Rolling Average Baseline for Total Traffic. (/day)               | 69 |
| 6.3  | One-Month Rolling Maximum Baseline for ICMP type 11                        | 69 |
| 6.4  | Price Channels from 445/tcp Dest. Port Packet Counts (Conficker.A) (/day). | 71 |
| 6.5  | Price Channels from Average Packet Size (Scan.A) (/day).                   | 72 |
| 6.6  | Bollinger Bands (%b) from 445/tcp Dest. Port Packet Counts (/day).         | 73 |
| 6.7  | Keltner Channels from 445/tcp Dest. Port Packet Counts (/day).             | 74 |
| 6.8  | Moving Average Envelopes from 445/tcp Dest. Port Packet Counts (/day).     | 74 |

# List of Tables

|      |  |    |
|------|--|----|
| 2.1  | Comparison of Network Telescope Size and Data Captured . . . . . | 13 |
| 2.2  | Results from CoralReef’s crl_stats Application . . . . .         | 17 |
| 3.1  | RU1 and RU2 Dataset Properties. . . . .                          | 26 |
| 3.2  | Descriptive Statistics for Destination IP. . . . .               | 31 |
| 3.3  | Top Destination IPs. . . . .                                     | 31 |
| 3.4  | Descriptive Statistics for Source IP. . . . .                    | 32 |
| 3.5  | Table of IP by Time Period and IP Class. . . . .                 | 33 |
| 3.6  | TTL Modes. . . . .   | 34 |
| 3.7  | Descriptive Statistics for TTL. . . . .                          | 34 |
| 3.8  | Descriptive Statistics for Packet Size. . . . .                  | 35 |
| 3.9  | Counts by Top Packet Sizes. . . . .                              | 35 |
| 3.10 | Descriptive Statistics for TCP Source Port Counts. . . . .       | 36 |
| 3.11 | Top TCP Source Ports. . . . .                                    | 36 |
| 3.12 | Descriptive Statistics for UDP Source Port Counts. . . . .       | 37 |
| 3.13 | Top UDP Source Ports. . . . .                                    | 37 |
| 3.14 | Descriptive Statistics for TCP Destination Ports. . . . .        | 38 |
| 3.15 | Dataset Coverage of TCP Destination Ports. . . . .               | 38 |
| 3.16 | Descriptive Statistics for UDP Destination Ports. . . . .        | 38 |
| 3.17 | Dataset Coverage of UDP Destination Ports. . . . .               | 39 |

---

|      |   |     |
|------|---|-----|
| 3.18 | Descriptive Statistics for ICMP Types. . . . .                | 39  |
| 3.19 | Dataset Coverage of ICMP Types. . . . .                       | 40  |
| 3.20 | Top OS Types. . . . .   | 40  |
| 3.21 | Packets with Related Characteristics. . . . .                 | 41  |
| 4.1  | Major DDoS Incidents Observed. . . . .                        | 46  |
| 4.2  | Emergence of Conficker Variants. . . . .                      | 47  |
| 4.3  | Cases of Net Send Malware. . . . .                            | 48  |
| 4.4  | Cases of Depth Scanning. . . . .                              | 49  |
| 4.5  | Cases of Breadth Scanning. . . . .                            | 49  |
| 5.1  | Incident Listing for Figure 5.2 (Total Packet Count). . . . . | 55  |
| 5.2  | Operating Systems Groupings. . . . .                          | 63  |
| 7.1  | Identification Ability of Count-Based Metrics. . . . .        | 78  |
| 7.2  | Identification Ability of Aggregated Metrics. . . . .         | 81  |
| 7.3  | Identification Ability of Grouping Metrics. . . . .           | 84  |
| 7.4  | Metrics Chosen for Identification Aids. . . . .               | 86  |
| 7.5  | Identification Ability of Baseline Analysis. . . . .          | 87  |
| 7.6  | Identification Ability of Technical Analysis. . . . .         | 90  |
| B.1  | Bogon IP Netblocks . . . . .                                  | 115 |
| B.2  | Commonly Used TCP Ports . . . . .                             | 115 |
| B.3  | Commonly Used UDP Ports . . . . .                             | 115 |
| B.4  | Commonly Used ICMP Types . . . . .                            | 116 |

*Fundamental characteristics of the Internet are perpetually challenging to research and analyze, and we must admit we know little about what keeps the system stable. As a result, researchers and policy makers currently analyze what is literally a trillion-dollar ecosystem [the Internet] essentially in the dark, and agencies charged with infrastructure protection have little situational awareness regarding global dynamics and operational threats.*

K. Claffy et. al, *Internet Mapping: from Art to Science*

# 1

## Introduction

THE Internet has created a rich environment for the formation of virtual communities (Preece, Maloney-Krichmar and Abras, 2003), electronic commerce (Scacchi, 1994), distributed educational (Kumar and Kush, 2006) and e-government utilities (OECD, 2003). These entities provide what is becoming an essential source of information and services to the global populace. Regrettably, there are a number of individuals who attempt to disrupt or subvert the functioning of these entities for their own nefarious purposes (BBC, 2010; IC<sup>3</sup>, 2009; Vijayan, 2011; Zetter, 2011). Internet-wide incidents such as the Morris worm (Orman, 2003), SQLSlammer (Ray, 2004) and Conficker (US-CERT, 2009a) illustrate the need to monitor the type of traffic traversing the Internet as each case demonstrates the effect that large scale incidents can have on the performance of the Internet.

Early identification of potential incidents such as the work done by Zou, Gong, Towsley and Gao (2005) places information security professionals in a position to ready organizations and individuals. Such action reduces the financial, administrative and technical impact of large scale incidents (Danchev, 2009; Lemos, 2003; Maurer, 2003). Researchers and network analysts have traditionally made use of logs from firewalls and intrusion

detection systems in an attempt to identify new strains of activity on the Internet (Dell Computer Corporation, 2009). The primary issue with this approach is the difficulty in separating legitimate traffic, that is traffic solicited by users, from malicious activity as the presence of legitimate traffic may influence incident identification techniques. To combat this issue researchers experimented with technologies that exclusively observe malicious traffic such as honeypots and network telescopes. While honeypots are a useful way of monitoring popular techniques employed by hackers (Mokube and Adams, 2007), this thesis focuses on the use of data obtained from network telescopes.

The network telescope approach to monitoring is to observe the traffic received from a portion of IP space that has no legitimate hosts. It follows that all traffic received must be unsolicited which neatly solves the issue of separating the legitimate from illegitimate traffic. Data obtained from network telescopes provides insight into the composition of malicious traffic traversing the Internet and can be used as an early indicator of potential outbreaks.

This introductory chapter defines the thesis problem statement, the research objectives, the research approaches, the research scope and then finally outlines the remaining chapters.

## 1.1 Problem Statement

A case in point of the real and virtual world damage that can be caused by malicious activity is exemplified by SQLSlammer. This worm indirectly disabled the IT components of the 911 emergency service in Washington and caused a general degradation in the quality of service of Internet traffic during late January 2003 (Lemos, 2003). This thesis considers the idea of network incidents, such as SQLSlammer, which are defined as large-scale points of interest where there is a considerable change in the makeup of malicious traffic transmitted across the Internet. The rampant spread of Conficker (Aben E., 2000; US-CERT, 2009a), SQLSlammer (CERT, 2003; Litchfield, 2010; Ray, 2004) and the Witty Worm (Moore and Shannon, 2004) are all excellent examples of network incidents.

Due to the modern day reliance of individuals and corporations on the Internet, it is important to be able to identify new attack vectors used by malware authors in order to effectively mitigate future threats. One popular way to predict potentially damaging attacks is to use traffic captured by network telescopes to act as a source of data from

which unusual activity can be identified.

However, there are a number of challenges that arise when attempting to make use of this data. Analysis of network telescope datasets is complicated by the large quantity of data present, the number of subdivisions within the data, the interplay of incidents within the data and the persistence of older exploits

Organizations such as the Co-operative Association for Internet Data Analysis (CAIDA)<sup>1</sup>, Team Cymru<sup>2</sup>, Arbor<sup>3</sup> and the Protected Repository for the Defence of Infrastructure Against Cyber Threats (PREDICIT)<sup>4</sup> have already made use of data collected by their network telescopes to keep abreast of current trends in malware. However, there is very little documented research on approaches to analysing network telescope data. This thesis aims to document a selection of said approaches to analysing network telescope datasets by evaluating the effectiveness of approaches and metrics to identify malicious activity.

## 1.2 Research Objectives

The research described in this thesis was constructed such that the following objectives are achieved:

- Identify incidents and points of interest in datasets to be used to evaluate the effectiveness of approaches to analysing network telescope datasets.
- Construct a number of metrics and evaluate their effectiveness to measure malicious activity on the Internet.
- Investigate methods of analysis used in existing fields of data analysis that would be appropriate for evaluating network telescope data.
- Evaluate the effectiveness of these new techniques by considering malicious activity observable in the datasets.

---

<sup>1</sup><http://www.caida.org/home/>

<sup>2</sup><http://www.team-cymru.org/>

<sup>3</sup><http://www.arbornetworks.com/>

<sup>4</sup><https://www.predict.org/>

## 1.3 Research Approach

The research discussed in this thesis was constructed as an introductory study on approaches to analysing network telescope datasets. A basis for this research was formed by considering previous approaches to network telescope data analysis. This was achieved by conducting a study of existing literature detailing the structure and usage of network telescopes. This was coupled together with a study of the CoralReef (Moore, Keys, Koga, Lagache and Claffy., 2001) network telescope data analysis framework and the techniques employed by Team Cymru.

Two datasets collected during previous research at Rhodes University were acquired and studied. The significant properties of the dataset were identified and investigated by performing descriptive statistical analysis. This analysis highlighted a number of points of interest which were then investigated by hand by parsing the packet traces using Wireshark to obtain a more comprehensive view. A number of incidents were identified through this investigation, including: Conficker, SQLSlammer, scanning and DDoS backscatter.

A number of metrics were outlined and investigated for effectiveness by applying the metrics to the datasets and observing the results. A number of approaches were drawn from existing research such as baseline analysis and technical analysis and were applied to network telescope datasets. The metrics and approaches considered were evaluated on the effectiveness in terms of their ability to identify incidents that had been found in the datasets and by considering the advantages and disadvantages that mentioned approaches provided.

## 1.4 Research Scope

While a large body of work is considered within this thesis, a narrowing of the domain is required to keep the length of this thesis manageable. The three major factors that affect the completeness of this study are the properties of the datasets, the number of datasets considered and the number of metrics and approaches considered.

This thesis analyses two datasets obtained from research performed at Rhodes University (Irwin, 2011). These two datasets, while small relative to the large datasets accumulated by CAIDA and Team Cymru, were chosen for two specific reasons. Firstly, network telescope datasets can be extremely large and as such are time consuming to process

and analyse. Initial attempts at processing a four GiB dataset often ran for periods of over 24 hours, though the processing time was improved by writing more efficient SQL statements and through better use of database indices. Secondly, it is difficult to obtain network telescope datasets as access is often restricted. Some organizations consider such data to be a competitive advantage and thus are unwilling to share data for comparative purposes.

There is no bound on the number of metrics that could be derived from data obtained from network telescope if it is considered that the time period to aggregate data from varies. Further, there are many ways to approach data analysis and numerous fields to investigate. It follows that this thesis will lack completeness in the approaches applied and this is accepted as the nature of this research approach. To re-iterate, this thesis is a first attempt to explore this topic of study.

## 1.5 Document Conventions

The conventions that are adhered to within the remainder of the body of this document are as follows:

- Where mention is made to an organization or application the URL for the associated website shall be noted at the bottom of the page as a footnote.
- Where values are quoted these shall be rounded to three decimal places. To deal with extremely large or small values scientific notation shall be used for brevity.
- The ‘Year-Day-Month Hour:Minute:Second’ form is used to format dates within the text. The format ‘Month-Abbreviated Day’ shall be used when dates are the units on the  $x$ -axis of graphs.
- Ports are listed according to the format {port number}/{protocol}. For example, the HTTP protocol port as assigned by the Internet Assigned Numbers Authority (IANA)<sup>5</sup> is listed as 80/tcp.
- CIDR (Classless Inter-Domain Routing) notation is used to describe sections of the IPv4 space.

---

<sup>5</sup><http://www.iana.org/>

- The shorthand notion ‘/x’ is used for ‘per x’ to reduce labels such as ‘number of packets per day’ to ‘number of packets /day’.
- A number of terms are abbreviated in tables and figures to ensure that these structures fit neatly within the confines of the page. These abbreviations are listed in brackets as follows: source IP (src. IP), destination IP (dst. IP), average (ave.), percentile (perc.), maximum (max) and minimum (min).
- Where specific mention is made to IPs belonging to the network telescopes used in this thesis, the first three octets shall be masked as a.b.c. This is to ensure that the network telescope is not targeted by hackers or ignored by worms and other malicious traffic in the future.
- All histograms are bucketed with  $n = 100$

## 1.6 Document Outline

The remainder of this thesis is organized into a number of chapters as follows below:

- Chapter 2 provides a summary of the relevant background required to understand the main thesis body. This includes a literature review of network security, network telescope fundamentals, cases studies and an overview of the data analysis techniques used in this thesis.
- Chapter 3 describes how the datasets were obtained, structured and summarized. A breakdown of the relevant variables is provided together with a descriptive statistical analysis of these variables.
- Chapter 4 provides background on major incidents that were observed in the dataset. The characteristic behaviour of the incidents identified are discussed together with suggested metrics that may be effective in their identification.
- Chapter 5 investigates a number of metrics that can be formed from network telescope data. This chapter deals with the issue of defining normality for an incident and explores a number of possible metrics constructed from counts, aggregation and grouping approaches.
- Chapter 6 considers techniques to aid in incident identification. Baseline analysis and technical analysis are the main techniques considered within this chapter.

- Chapter 7 summarizes the results obtained from the approaches discussed in Chapter 5 and Chapter 6 as applied to incidents noted in Chapter 4.
- Chapter 8 concludes the thesis. The outcomes of the research are considered together with potential future work in the field of network telescope data analysis.

*Intelligence analysts should be self-conscious about their reasoning processes. They should think about how they make judgments and reach conclusions, not just about the judgments and conclusions themselves.*

Richard J. Heuer former CIA operative, Psychology of Intelligence Analysis

# 2

## Literature Review

THE core topics of this thesis, data analysis and network telescope datasets, are extensive technical fields that require background to familiarize the reader adequately. While the analysis of network telescope data is a relatively new field of inquiry there is some existing research that can be drawn upon to serve as a basis. Data analysis is a broad subject and has been applied to a variety of disciplines including economics, computational physics and engineering processes. Inspiration can be derived from these fields to formulate approaches to analysing network telescope datasets. This chapter is organized as follows below:

- Section 2.1 briefly describes networking concepts relevant to this thesis. Topics covered include: the growth and importance of the Internet, the motivation for network security and the types of network traffic observed on the Internet.
- Section 2.2 provides background information on the components, properties and limitations of network telescopes.
- Section 2.3 is a case study of the techniques employed by Team Cymru and CAIDA to analyse the state of malicious activity on the Internet.

- Section 2.4 describes the approaches investigated from other fields as applied in Chapters 5 and 6. Fields explored include: security metrics, moving averages, technical analysis, baseline analysis and descriptive statistics.
- Section 2.5 summarizes the concepts covered in this chapter.

## 2.1 Networking Concepts and Internet Monitoring

From humble beginning as the interconnection of four prominent American universities forming the Advanced Research Projects Agency Network (ARPANET) (Leiner, 2003), the Internet has grown vastly in geographical coverage, physical infrastructure and available content. This is illustrated by the fact that in the last sixteen years global Internet penetration has increased from a meagre 0.4% in December 1995 to 30.9% in June 2011 (Internet World Stats, 2011). As a result of this increased availability, the Internet has become a part of the lives of many individuals and an integral component to the functioning of organizations (Fuchs, 2007). It follows that ensuring the continued optimal functioning of Internet should be of great importance to governments and other organizations.

Network worms have shown the devastating effect that malicious code can have on the functioning Internet with little forewarning. The Morris worm, the first major Internet worm, illustrates this point well having infected 10% of the Internet in 1989 causing a significant slow-down of the Internet as data lines became saturated (Orman, 2003). Fast spreading worms such as SQLSlammer illustrate the rapid rate at which these infections can grow as shown by the fact that approximately 90% of hosts susceptible to SQLSlammer were infected in 10 minutes (Ray, 2004). These incidents make a strong case for monitoring the state of malicious activity on the Internet.

### 2.1.1 Types of Network Traffic

A wide variety of applications, services and protocols have been developed for the Internet resulting in an assortment of traffic types observable by hosts connected to the Internet. This section differentiates traffic based upon the datagram type used, the legitimacy of the traffic and the active or passive nature of the traffic.

The datagram type chosen by a particular application is usually indicative of the intention of the traffic. The majority of traffic found travelling the Internet will usually make use of

either TCP (Transmission Control Protocol), UDP (User Datagram Protocol) or ICMP (Internet Control Message Protocol). Applications that require end-to-end reliability will tend to use TCP as this datagram type supports the retransmission of lost data (Information Sciences Institute, 1981b). UDP is often used by applications that do not require this transmission guarantee but instead prefer a smaller packet header (Postel, 1980). ICMP is a protocol which functions as a messaging system to monitor and manage the state of IP enabled devices (Postel, 1981).

An alternative classification for network traffic may be formed by considering whether traffic is legitimate, anomalous or malicious. Legitimate traffic is defined as traffic that was generated as a request or as a solicited response to a specific user or application request that reaches its destination. For example, the traffic generated from a user requesting a web-page and the response to that request are considered to be legitimate. Anomalous traffic is defined as traffic that was generated and received by a host but is unexpected in its content or destination. This could occur due to software misconfiguration or hardware failure. Traffic may be defined as being malicious if its intention is to cause damage or acquire intelligence. DDoS, ping floods and scanning are all examples of malicious traffic (Irwin, 2011).

Traffic may be broadly classified as passive or active. Passive traffic does not intend to evoke a response from the destination host. TCP packets with the RST flag set and ICMP Echo replies are examples of passive traffic. On the other hand, active traffic is intended to evoke a response from the destination host. Examples of active traffic are ICMP pings or TCP packets with legitimate combinations of the SYN, ACK, FIN or URG flags (Irwin, 2011).

Being aware of the types of malicious attacks and the reasoning behind such attacks is important when designing an effective monitoring system as it allows analysts to construct their systems according to specific parameters. The next section of this thesis provides a brief overview of malicious activity on the Internet and the need for network monitoring.

### 2.1.2 Motivation for Network Security

Malicious activity on the Internet is an ever evolving field of study with attacks such as Cross Site Scripting, SQL injection, DDoS, worms, viruses, flooding, smurfing and identity theft. This section describes some of the malicious activity conducted on the Internet which serves to suggest the need for research in network security.

Financial remuneration is a popular reason that hackers target individuals and organizations. This is the postulated reason behind the botnet infrastructure developed by the creator of the Conficker worm, to act as a network of ‘zombies’ that can be paid to target a particular host for DDoS (Nahorney, 2009). Other examples of this rationale being employed include a wide variety of key loggers and spyware applications that have been used to collect information on users which can then be sold in bulk.

Political and socio-economic beliefs of a group can result in online organized protest against an organization, group or individual. Members of the anonymous image board 4chan<sup>1</sup> perpetrated an organized DDoS attacks using a SYN flood DoS tool known as Low Orbit Ion Cannon (LOIC)<sup>2</sup> against the MPAA (Motion Picture Association of America)<sup>3</sup>, RIAA (Recording Industry Association of America)<sup>4</sup>, Aiplex<sup>5</sup>, Sony<sup>6</sup> and other organizations (Anderson, 2010; Corrons, 2010; Reuters, 2010; Tsotsis, 2010).

Worms and viruses can cause a significant amount of damage to the functioning of networks. An example of this is SQLSlammer (Ray, 2004), a worm that exploited two buffer overflows in Microsoft SQL Server Desktop Environment 2000 (MSDE) infecting an estimated 75000 hosts within a few minutes after initial outbreak. SQLSlammer flooded portions of the Internet resulting in indirect DDoS as routing equipment could not deal with the sudden and sharp increase in traffic. Perhaps the most interesting property of SQLSlammer is the fact that it is still present on the Internet to this day, seven years after initial outbreak. This is particularly surprising given that the worm can be removed simply by restarting the infected machine as the worm is memory resident.

The cases noted above highlight the importance of safeguarding information technology infrastructure at both an organizational and individual level. Thus it follows that there is a need for Internet monitoring and related research to understand the nature of attacks and to provide warning as new attacks emerge as this allows for countermeasures to limit the spread of damage caused by such an incident.

The next section of this chapter details the structure and related properties of network telescopes while considering some of the limitations that this type of monitor may experience.

---

<sup>1</sup><http://www.4chan.org/>

<sup>2</sup><http://sourceforge.net/projects/loic/>

<sup>3</sup><http://www.mpaa.org/>

<sup>4</sup><http://www.riaa.com/>

<sup>5</sup><http://www.aiplex.com/>

<sup>6</sup><http://www.sony.com/>

## 2.2 Network Telescope Concepts

Network telescopes provide a looking glass into the world of malicious activity on the Internet (Harder, Johnson, Bradley and Knottenbelt, 2004). This is accomplished by listening to the traffic bound to an address space on which there are no hosts present (Moore, 2002). The generic structure of a network telescope involves a device capable of IP routing and packet capture, a connection to the Internet and a free block of IP space (Baker, Harrop, and Armitage, 2010) as illustrated in Figure 2.1.

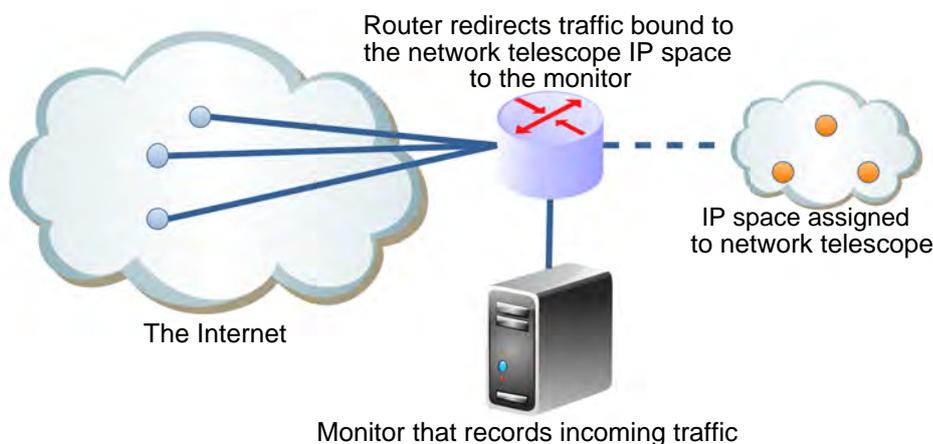


Figure 2.1: Components of a Network Telescope

It can be assumed that all traffic observed from a network telescope is malicious or anomalous from the fact that the traffic is completely unsolicited (Harrop and Armitage, 2005; Moore, Shannon, Voelker and Savage, 2004). The possibility of new Internet incidents may be inferred by measuring the amount of traffic received according to specific criteria. The properties of a network telescope can greatly affect whether certain incidents are observable and how clear these incidents are in network telescope datasets, as is described in the following subsection.

### 2.2.1 Properties of Network Telescopes

The three major properties to be considered when discussing network telescopes are the size of the address block monitored, the location of the address block and the configuration of the device used to record traffic received.

The size of the address block increases the probability that a network telescope will receive a packet from a malicious source assuming that all sources of malicious traffic pick their destination hosts at random. This probability can be expressed as shown in equation (2.1) where  $n$  is the number of hosts assigned to the network telescope space. It is approximate as not all portions of the IPv4 space are routable as some are reserved for other purposes as shown in Addendum B.

$$P(\text{packet}) \approx \frac{n}{256^4} \quad (2.1)$$

To illustrate this point a comparison is made between the number of packets captured in a month period by one of CAIDA’s network telescopes and data captured from a network telescope used in this thesis (RU1) for the same time period of May 2010. Additional information is provided on RU1 in Chapter 3.

| Dataset Name | Netblock Size | P(Random Packet)               | Data (GiB/s) |
|--------------|---------------|--------------------------------|--------------|
| CAIDA        | $\approx /8$  | $\approx 3.906 \times 10^{-2}$ | 562.5        |
| RU1          | $/24$         | $3.960 \times 10^{-7}$         | 1.71         |

Table 2.1: Comparison of Network Telescope Size and Data Captured

It is clear from Table 2.1 that the size of the network telescope has a major effect on the amount of data received from both theoretical and experimental standpoints. Whether this affects the ability of smaller telescopes to observe major incidents is an important question. Work conducted by Harder, Johnson, Bradley and Knottenbelt (2004) suggests that while a smaller telescope receives a lesser sampling of the total anomalous traffic on the Internet, there is still sufficient data to detect major changes in the composition of malicious traffic.

The logical location of a network telescope can play an important part in what activity is observable by a network telescope. A targeted attack made against a specific country illustrates this point as this attack would be unobservable by network telescopes outside of this IP space (Beck, Festor and State, 2007; Moore, Shannon, Voelker and Savage, 2004).

The configuration of a network telescope can have a major effect on what incidents can be observed as certain configurations will reduce the observable space but may offer

specific advantages. Passive network telescopes are configured specifically to only receive traffic and never reply to any traffic. As a result protocols that require some form of acknowledgement or handshake will be unable to form a connection. For example, it is not possible to form a TCP session with a passive network telescope as the TCP three-way handshake can never complete (Information Sciences Institute, 1981b). As the TCP handshake cannot complete no data transfer can occur and thus the payloads of TCP packets are unavailable to passive network telescopes. It is noted that specific configuration choices may be made to reduce potential for attacks to be made against the network telescope due to the fact that such monitors are often targeted by hackers (Bethencourt, Low; Shinoda, Ikai and Itoh, 2005).

Having listed the composition and core properties of a network telescope it is now important to consider the limits on what is observable by a network telescope as is discussed in the following section.

### 2.2.2 Limitations of Network Telescopes

It is critical to note the type of incidents that are observable by network telescopes are limited by a number of factors. For example, a worm's propagation algorithm could be designed such that it will avoid attempting to scan or infect IPs that belong to known network telescopes. It is for this reason that researchers will not reveal the IP space that their telescopes monitor as this may negatively impact the sample of malicious traffic observed. Three cases where it is unlikely that a network telescope will observe any unusual activity are provided below:

- Cross Site Scripting (XSS), Cross-site request forgery (CSRF), SSL man-in-the-middles and other web based attacks that target users through their browsers are unlikely to be observed by a network telescope. These attacks exploit vulnerabilities in browsers and the naivety of users through specifically crafted web pages. A system that monitors this sort of malicious activity would need to proactively scan for such pages.
- Worms, viruses and trojans that use e-mail or infected files as their sole infection vector are unlikely to be observed by network telescopes. E-mail as a vector is almost certainly going to be unobservable as the algorithms used to target victims generate e-mail addresses and not specific IPs. Infected files require that the victims

download the infected file to their machine for infection to take place. To observe this type of behaviour a monitor would have to be actively scanning for these files.

- Targeted SYN flooding, ping flooding and other approaches to DDoS may be unobservable by network telescopes assuming that the attacker is not spoofing their IP space or if none of the spoofed IPs fall in the IP space of a network telescope.

## 2.3 Case Studies

To understand previous approaches to analysing network telescopes it is necessary to study existing systems. This section takes a brief look at how CAIDA and Team Cymru employ network telescopes to monitor the state of malicious activity on the Internet.

### 2.3.1 CAIDA

The Cooperative Association for Internet Data Analysis (CAIDA) is a partnership between the University of California San Diego and the information security industry for the purpose of Internet related research. CAIDA collects large quantities of data for analysis and visualization in order to gain a better grasp of the malicious activity occurring on the Internet. CAIDA's particular research areas include the analysis of network topologies, traffic analysis, malicious activity analysis and the visualization of Internet traffic (Claffy, Hyun, Keys, Fomenkov and Krioukov, 2009).

CoralReef<sup>7</sup> is a unified set of applications developed by CAIDA for the purpose of constructing reports from real-time or pre-recorded network traffic for data analysis (Keys, Moore, Koga, Lagache and Claffy, 2001). Many of the real-time monitors employed by CAIDA make use of CoralReef as a way to communicate the latest observed data via HTTP (Hyun, 2008). Figure 2.2 shows data captured from from CAIDA's Chicago monitor A passive network telescope (CAIDA, 2010b).

- The number of (unique) IPv4 packets received.
- The number of (unique) destination and source ports received.
- The number of (unique) ICMP types received.

---

<sup>7</sup><http://www.caida.org/tools/measurement/coralreef/description.xml>



## Passive Network Monitors

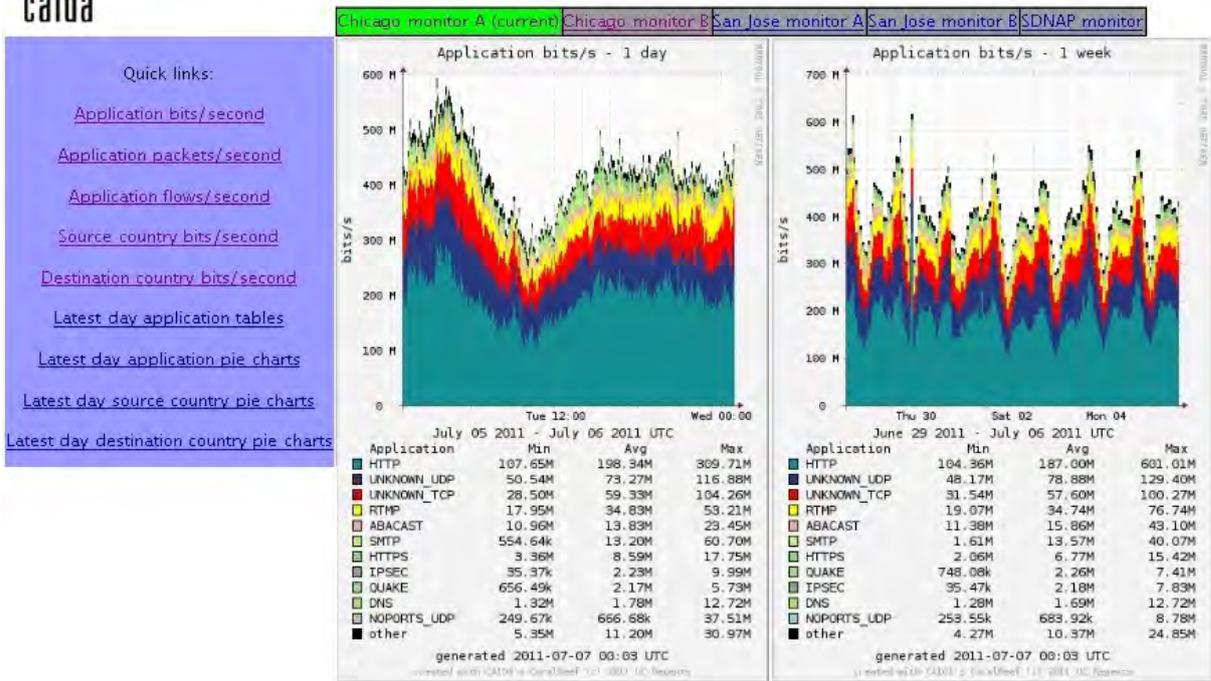


Figure 2.2: CoralReef Monitoring System (Chicago B Monitor)

| Property                               | Value                |
|--|----------------------|
| Maximum capture length for interface 0 | variable length      |
| Maximum capture length for interface 1 | variable length      |
| First timestamp                        | 1105660500.000004262 |
| Last timestamp                         | 1105660799.999996483 |
| Unique IPv4 addresses                  | 1406173              |
| Unique src IPv4 addresses              | 485775               |
| Unique dst IPv4 addresses              | 1026819              |
| Unique src ports                       | 64927                |
| Unique dst ports                       | 65536                |
| Unique ICMP type/codes                 | 0                    |
| Unknown encapsulation                  | 972                  |
| non-IP protocols                       | 0                    |
| non-IP pkts                            | 0                    |
| IPv4 pkts                              | 102830571            |
| IPv4 bytes                             | 62747963706          |
| IPv6 pkts                              | 6943                 |
| IPv6 bytes                             | 704205               |
| Flows                                  | 5402401              |

Table 2.2: Results from CoralReef's `crl_stats` Application

- The source country of the packet.
- The destination country of the packet.
- The number of (unique) source IP of the packet.
- The transport protocol type.

Using these variables analysts from CAIDA make inferences on the state of malicious activity on the Internet and provide data from which to study the nature of large scale worms. CAIDA researchers have observed the activity of major worms such as Conficker, SQLSlammer, Code-Red and the Witty through data obtained from their network telescopes (Moore, Paxson, Savage, Shannon, Stanifor and Weaver, 2003; Moore and Shannon, 2004; Moore, Shannon and Claffy, 2002).

### 2.3.2 Team Cymru

Team Cymru is a non-profit organization with the purpose of “Making the Internet a more secure place by monitoring malicious activity on the Internet” (Team Cymru, 2011b). Team Cymru provides valuable Internet services including a malware hash registry, an Autonomous System (AS) mapping service and the Internet Garbage Meter. Team Cymru has contributed to the field of information security research through the design of network telescopes and the monitoring of malevolent activity on the Internet.

## 2.4 Data Analysis Techniques and Approaches

Data analysis is a process that transforms raw data into a state from which useful information can be extracted and conclusions may be drawn. Example transforms include: removing outliers, normalization through logarithms or grouping data into categories (Nison, 2011). Data analysis is a well established field of study with a number of generic approaches such as linear regression and field specific techniques such as spectral analysis. This section provides background on security measures and metrics, baselines, moving averages, technical analysis and descriptive statistics.

### 2.4.1 Security Measurements and Metrics

A security measurement, according to Payne (2006) provides a “Single-point-in-time-view of specific and discrete factors which gives insight into the state of information security within an organization”. The average strength of user passwords for a particular month is an example of a security measure that can be used in conjunction with other measures to estimate the risk of illegitimate access to an organization’s resources.

Security metrics are created by comparing measures taken under differing conditions and drawing conclusions from the analysis of the values. The differing condition could be different time periods or different business units within an organization. Comparing the time taken to apply critical patches may be used as a metric to infer potential risk from worms and viruses (Jaquith, 2007).

It is a commonly accepted principle that to be able to effectively manage an activity it is critical to obtain accurate measurements of the properties related to the activity (Jaquith, 2007; Payne, 2006). Within the sphere of information security, security measures provide a way to measure risk and thus prioritize the actions taken to remediate potential danger (Payne, 2006).

### 2.4.2 Baselines

A common approach to measuring the change in a quantity is to construct a baseline from which comparison can be made. That is to say a baseline is a set of measurements taken at a particular instance or aggregated value over a period from which new readings can be compared against. Baselines allow for the establishment of expectations for a given quantity and thus provide a mechanism to identify unusual activity (Holman, Johnson and Bradley, 2009).

### 2.4.3 Moving Averages

Moving averages are a type of filter that aggregates data points in a time-specific window. The ‘window’ of a moving average is defined as the period from which values are incorporated as part of the calculation (Press, Teukolsky, Vetterling and Flannery, 1992). For example, a seven day window will include the values of the last seven days and

then calculate an average using the data obtained at each day. This allows an observer to consider what sort of behaviour is normal for a given window period. A number of moving averages exist such as the simple moving average, simple median average, exponential average, the adaptive moving average and Kruskal's moving average (Faires and Burden, 2005).

In this thesis two types of moving averages are considered: the simple moving average (SMA) and the exponential moving average (EMA). The simple moving average (SMA) provides an efficient and easy technique for calculating a moving average. Effectively the SMA is a windowed geometric mean, as shown by equation (2.2) (Faires and Burden, 2005).

$$SMA(x, n) = \frac{x_1 + x_2 + x_3 + \dots + x_n}{n} \quad (2.2)$$

A more complex version of the moving average can be formed by weighting the values in a window differently. The EMA is an example of such an approach in that the current value is considered the most informative value for the prediction of the next value and thus given the greatest weighting while the weight of older observation decays exponentially. Mathematically the EMA can be expressed as shown in equation (2.3) (Kirkpatrick, 2002).

$$EMA(x, n) = \frac{2x_n - x_{n-1}}{n + 1} \quad (2.3)$$

#### 2.4.4 Technical Analysis

John Bollinger, a prominent market analyst and the inventor of the Bollinger Band, defines technical analysis (TA) as “The study of market-related data as an aid to investment decision making” (Bollinger, 2002a). Technical analysis employs indicators, techniques and trading rules that consider the current and historical direction of prices together with other market related quantities in order to make better decisions with regards to the buying and selling of stocks. As quantities tend to fluctuate unpredictably in the short-term, moving averages and other filtering techniques are used to smooth out spikes in the dataset.

For some more introductory reading in economic theory and basic technical analysis the interested reader is recommended to read “Technical Analysis of Financial Markets” by (Murphy, 1999a).

### **Fundamental Assumptions**

In order to apply approaches from technical analysis it is important to consider some of the assumptions used in the field as many indicators require these statements to be true in order to be useful. The fundamental assumptions of technical analysis are as stated in (Edwards and Magee, 2001):

- Stock prices are determined solely by the interaction of demand and supply.
- Stock prices tend to move in trends.
- Shifts in demand and supply can be detected in charts.
- Chart patterns tend to repeat themselves.

The question as to whether these assumptions hold for the network telescope monitoring context is discussed in Chapter 7. Having stated the fundamental assumptions of technical analysis, the next section defines some of the technical analysis concepts discussed in this thesis.

### **Technical Analysis Concepts**

Some of the basic concepts of technical analysis are discussed in this section such as financial quantities, indicators, overbought, oversold, trend, crossovers, support, resistance, loop back periods and trading rules. Figure 2.3 is a stock chart of the SPY (Standard & Poor’s Depository Receipts) sourced from Yahoo! Finance<sup>8</sup> that has been annotated with the concepts that follow:

- Financial quantities are measures derived from objects such as the price of a stock, the value of financial indexes such as the NASDAQ (National Association of Securities Dealers Automated Quotations)<sup>9</sup> or the number of stocks that have been traded

---

<sup>8</sup><http://finance.yahoo.com/>

<sup>9</sup><http://www.nasdaq.com/>

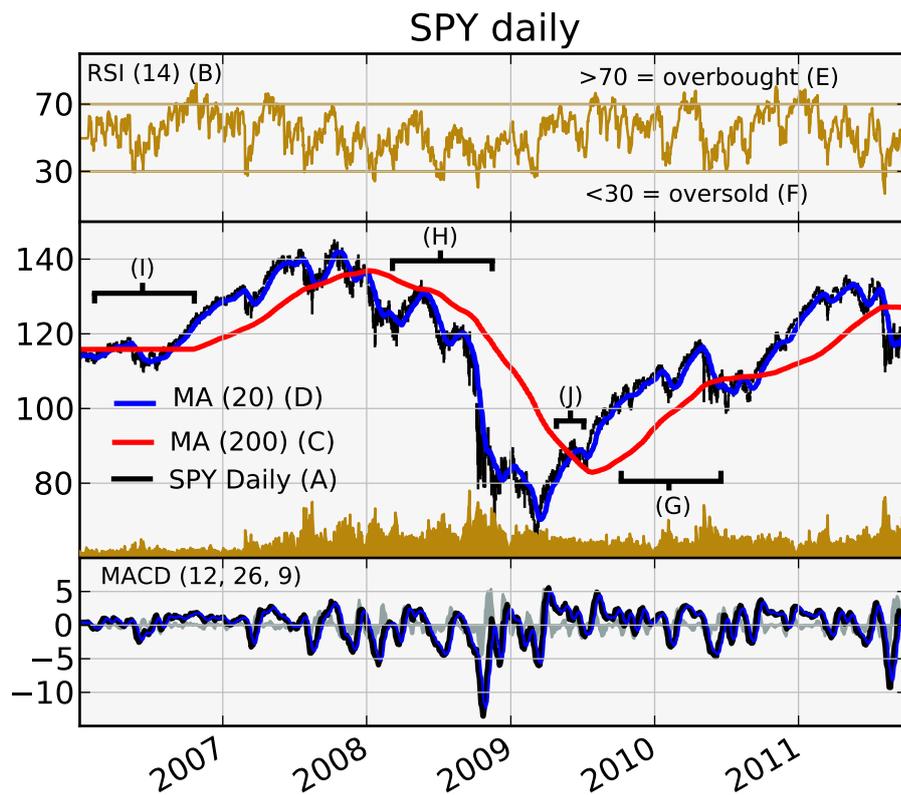


Figure 2.3: Financial Stockchart of SPY with TA Indicators.

in a particular day. These quantities are used on their own or as part of calculation to determine whether a particular action, such as buying or selling a stock, is favourable at a given time (Murphy, 1999a). In Figure 2.3 this is represented by the SPY daily as shown by point (A).

- Indicators are computations constructed from quantities such as the current price of a certain stock, the current value of a financial index or other financial values. Oscillators are a specialization of indicators that are limited to a specific range (Colby, 2002). The Relative Strength Index (RSI), a popular technical analysis indicator, as shown in Figure 2.3 by point (B) is an oscillator as it is limited to a value in the range  $[0,100]$ .
- A loopback period is defined as the last  $n$  periods that are to be considered as data for calculating technical analysis indicators (Murphy, 1999a). The MA(200), as shown by point (C) in Figure 2.3, is a 200 day moving average meaning that uses the last 200 days of data to calculate a single aggregated value.
- A quantity becomes overbought when it has risen too rapidly in the recent past

and is likely to suffer price decline in the short-term, this describes a situation where selling is likely to have a good outcome. Similarly, a quantity is said to be oversold when a quantities value falls rapidly usually implying that the quantity is undervalued which is a signal to buy said stock (Edwards and Magee, 2001). It is important to note that the concepts of overbought and oversold are defined for a specific indicator. For example, RSI defines overbought when the RSI value exceeds 70 as shown by point (E) and oversold when the RSI is below 30 as shown by point (F).

- Trend, in the technical analysis context, may loosely be defined as the general direction that a quantity is moving. A quantity is said to be experiencing an upward trend when the quantity is increasing overall as shown by point (G), a downward trend is when the quantity is decreasing overall as shown by point (H) and a sideward trend occurs when the quantity experiences no significant movement as shown by point (I) (Edwards and Magee, 2001).
- A crossover occurs when there is an intersection between two indicators or an indicator and a quantity. When a crossover occurs it is thought that either the quantity has become oversold or overbought. Crossovers are often used as justification to suggest that a change in the trend has occurred. This can be used as a signal in a trading system to prompt the buying or selling of stocks as appropriate (Edwards and Magee, 2001). Point (J) in Figure 2.3 is an example of crossover.
- Support represents the values from which a decreasing quantity struggles to descend (Murphy, 1999b). Resistance is the opposite of this and represents the values from which the quantity struggles to continue to rise (Murphy, 1999b). It is difficult to show support and resistance in a single image and as such this topic is discussed in more detail as needed in Section 6.2.
- Trading rules are triggers that traders use to determine when to buy or sell stocks. A common approach to this is to construct an indicator with a specific loopback, such as the fourteen day SMA for a financial index, then wait for a crossover to occur and sell as appropriate (Edwards and Magee, 2001).

Having discussed some of the core concepts of technical analysis the next section discusses trading bands, a popular techniques used to identify significant market behaviour.

## Trading Bands

Trading bands are used for a variety of different trading approaches within economics. A trading band is simply an envelope or a channel around a quantity that has been observed for a number of periods. These bands are constructed by plotting two lines that are a distance away from a quantity or a measure of central tendency for that quantity. Trading bands define an upper band which is generally above the quantity and a lower band which is generally below the quantity. The distance of these lines is dependent on the trading band type but is usually either a measure of volatility, a fixed distance or a function of some other related value (Bollinger, 2002b). Examples of trading bands include: Donchian price channels, Bollinger bands and Bollinger indicators, Keltner channels and moving average envelopes. These techniques are applied to network telescope data in Chapter 6.

### 2.4.5 Descriptive Statistics

Descriptive statistics allow for the reduction of massive datasets into a simple and summarized form that provides insight into the nature of the population being studied. Descriptive statistics are often used as a ‘first-look’ into a dataset, laying the foundation for more complex analysis. Descriptive statistics tends to focus on the distribution of data particularly the central tendency, dispersion of values and the association between variables (Coursey, 2003).

A common way to perform descriptive statistics is to calculate the so called ‘number summaries’ for the data which includes the 25<sup>th</sup> percentile, 75<sup>th</sup> percentile, mean, median, maximum, minimum and the standard deviation for the datasets together with histograms to show how the data has been distributed. This is the approach taken in Section 3.4 of this thesis.

## 2.5 Summary

This chapter has provided background on the core topics discussed in this thesis. Section 2.1 began by highlighting the growth of the Internet and its importance to modern society. Definitions for malicious, anomalous and legitimate traffic were provided as well. The section continues by discussing the malicious activity on the Internet and suggests the need to monitor this activity.

Section 2.2 introduces network telescopes as a way to monitor the malicious activity on the Internet. A network telescope was defined as having a IP device capable of recording network traffic bound to an IP space for which there are no legitimate hosts. Section 2.2.1 notes the important properties of a network telescopes. The size, logical location of the telescope and the configuration of the recording device were identified as significant network telescope properties. In particular the effect the size can have on the amount of malicious traffic is investigated.

In order to gain a sense for current approaches to network telescope data analysis a case study was performed in Section 2.3. Two organizations, Team Cymru and CAIDA, were briefly discussed with regards to their approach to research and analysis of network telescope data. CAIDA consider a number of variables such as the number of packets received per source country and bytes received per protocol from their /8 network telescope. Team Cymru make use of a number of smaller network telescopes which they aggregate to a so-called ‘Internet Garbage Meter’ to measure changes in the amount of malicious traffic received by nodes on the Internet.

The final section of this chapter, Section 2.4, details the data analysis approaches used in future chapters. Security metrics were defined together with a motivation for their usage. Simple and exponential moving averages were discussed and defined in Section 2.4.3. A brief introduction to technical analysis and the related terminology and concepts is provided in Section 2.4.4. This section is concluded by outlining an approach to descriptive statistics through the use of numerical summaries and histograms.

Having considered network telescopes and a number of possible data analysis approaches it is important to understand the nature and structure of network telescope datasets before these approaches can be applied in a sensible manner. With this in mind, Chapter 3 details the origin and properties of the datasets studied together with a descriptive statistics study of the significant variables.

*Data is a precious thing and will last longer than the systems themselves.*

Sir. Tim Berners-Lee

# 3

## Dataset Overview

**T**HIS chapter describes the nature and organization of the datasets that were used for analysis and testing in the remainder of this thesis. It is necessary to discuss the origin of the datasets used in this thesis in order to conduct any meaningful research on behaviour observed as this provides context to the observations made and may explain dissimilarities observed in other datasets. Further, understanding the properties and distribution of the datasets aids in the selection of metrics and identification techniques as this may suggest which approaches are unlikely to be appropriate. The remains of this chapter are structured into sections as listed below:

- Section 3.1 details the origin and properties of the network telescope datasets used in this thesis.
- Section 3.2 extracts the significant variables from the raw data. The function of each variable with regards to networking and importance of each variable to incident identification is also provided.
- Section 3.3 outlines the process of summarizing the datasets in order to reduce the amount of processing required to analyse the datasets.

- Section 3.4 provides an elementary statistical study of the significant variables identified in Section 3.2.
- Section 3.5 summarizes the issues addressed in this chapter.

## 3.1 Dataset Origins

The datasets studied in this thesis were obtained from network telescopes deployed by Dr. Barry Irwin as part of his PHD thesis (Irwin, 2011). Irwin saw the need for a network telescope on the African continent as at the time no known telescope existed within the AfriNIC<sup>1</sup> Internet Routing Registry IP block. A network telescope hosted in South Africa offered two significant advantages over analysing packet captures from other organizations. Firstly, access to the data would be significantly simpler as obtaining data from other network telescopes can be administratively challenging, as many organizations are unwilling to provide complete access to their data and logistically challenging as the datasets tend to be very large. Secondly, a telescope placed in IP space administered by AfriNIC may shed a different perspective on trends in malicious traffic as compared to telescopes in other IP spaces allowing for a contrast.

The datasets used in this research were obtained from two passive network telescopes with a total of 256 IP addresses per network telescope. The particular details relating to the configuration and setup of these particular telescope are detailed in (Irwin, 2011). RU1 was deployed outside of Rhodes University due to the security concerns expressed by the IT division of Rhodes University. Deployment of RU2 followed later but was designated an IP space belonging to Rhodes University. A summary of the significant properties of these datasets is provided in Table 3.1, showing that RU1 is an order of magnitude larger than RU2 with regards to the number of packets received due to the longer period of deployment.

| <b>Properties</b> | <b>RU1</b>               | <b>RU2</b>               |
|-------------------|--------------------------|--------------------------|
| Period            | 2005-08-03 to 2009-09-30 | 2009-08-20 to 2010-08-31 |
| Allocated Network | TENET                    | Rhodes University        |
| No. Packets       | 40801854                 | 2909634                  |

Table 3.1: RU1 and RU2 Dataset Properties.

---

<sup>1</sup><http://www.afrinic.net/>

Within the context of this thesis, RU1 is used as data for the analysis techniques tested in Chapter 5 through to Chapter 7. The reasoning for this is that RU1 provides a good case study of the outbreak of Conficker as discussed in Chapter 4 and further it covers a far longer period allowing for a better understanding of metrics and techniques applied. RU2 is used in this section as a dataset to compare the statistical properties of the variables studied in RU1 and as such only numerical summaries of RU2 are considered in this chapter.

## 3.2 Decomposing Network Traffic

Network traffic presents a complex modelling scenario as each packet may be encapsulated in multiple layers with each layer having numerous fields. In the interest of reducing the analysis to manageable levels only components that are significant to incident identification are to be considered for analysis. To achieve this each packet will be thought to consist of a date-time stamp to indicate when the packet was received, a source and destination IP, a packet length, a time-to-live value (TTL), an inferred OS type and a datagram type. The datagram types to be considered in this simplified model are TCP, UDP and ICMP. Both UDP and TCP have a source and destination port while ICMP has a ICMP type property. Each of these properties is now discussed in more detail:

- **Date-Time Stamp:** The date-time stamp is a POSIX (Lewine, 1991) time-stamp with the time zone set to Greenwich Mean Time (GMT) +2. The date-time stamp itself is useful as the independent variable to measure other quantities against.

The date-time stamp can be used to detect sudden packet surges which may be indicative of worm based attacks, DDoS or port scanning.

- **Source IP Address:** The source IP address is represented by a 32-bit value implying a possible of  $2^{32}$  unique hosts (Information Sciences Institute, 1981a) and is used to identify the machine sending the packet. The actual space is considerably smaller than  $2^{32}$  hosts due to large chunks of the IPv4 space being unroutable. Further, portions of the routable IPv4 generate no malicious activity as not all hosts are infected or used by malicious individuals.

The source IP address provides information on how distributed a particular attack may be amongst the nodes of the Internet. If a single host is generating the majority of new malicious traffic it is likely to be a targeted attack such as scanning or DDoS.

Alternatively, observing traffic from a wide variety of hosts may imply a large scale worm attack. A factor worth considering is the so called ‘bogon’ IPs, these are IPs that should not be observed by the public Internet as they are reserved for specific functions such as internal networking and multicast addresses (IANA, 2011; Team Cymru, 2011a). A list of bogon IP space is provided in Appendice B.1.

- **Destination IP Address:** While the destination IP address is also a 32-bit value, the number of possible hosts is far smaller than for the source IP as a network telescope only represents a specific subsection of the Internet. Typically speaking a small network telescope will use a /24 or /16 subnet though larger telescopes do exist (CAIDA, 2010a).

Points of interest may occur when a specific IP or a specific section of telescope space receives more traffic than other hosts which may occur due to targeted scanning.

- **Time-To-Live (TTL):** The TTL is a 8-bit value that is used to prevent loops from occurring in networks by acting as a way to keep track of the number of hops made by a packet as it moves from node to node. When the TTL value reaches zero the packet is no longer forwarded (Information Sciences Institute, 1981a).

The TTL value is one of the factors that is used to fingerprint the OS type. It is also possible to infer the logical distance between the source and host from the TTL value.

- **Transport Protocol Type:** The three main transport protocol types considered in this thesis are the standard Transport Control Protocol (TCP), Universal Datagram Protocol (UDP) and the Internet Control Message Protocol (ICMP) protocols of the Internet Protocol Suite. While there are a number of other transport protocols, such as the Resource Reservation Protocol (RSVP) or Datagram Congestion Control Protocol (DCCP), these protocols are far less popular than the three main protocols.

The protocol type may be useful when identifying major shifts in the malicious traffic on the Internet. This would be true for worms such as Conficker, which caused a significant increase in the amount of TCP traffic observed, due to the high volumes of traffic that a worm of this nature generates.

- **Packet Size:** The packet size property for an IPv4 packet is a 16-bit value which implies a range of [0,65336] bytes of payload (Information Sciences Institute, 1981a). However, to improve link efficiency as a single large packet may congest a transmission link a Maximum Transmission Unit (MTU) is enforced by a network. The IPv4

protocol will pad all packets to a minimum of 62 bytes and the Ethernet protocol only allows a 1500 byte MTU.

If a network telescope receives a number of packets all of the same size for a size that has historically had low traffic it is possible to infer that there may be a new trend in malicious activity. Packets containing a worm payload will tend to be similar in size unless there is some form of random padding for example.

- **Source and Destination Port:** A 16-bit value denoting the port from which the packet originated from or is targeting.

As a port is usually associated with a specific application the targeted port is of interest as a sudden increase in the amount of traffic received by a specific port that previously has not been particularly active may allude to a new vulnerability.

- **ICMP Type:** As the name suggests it specifies what sort of ICMP message is contained within the packet. A listing of relevant ICMP types can be found in Addendum B.

An increase in a specific ICMP type may be indicative of a particular type of incident such as ping flooding or DDoS.

- **Operating System (OS):** can be inferred through the use of p0f<sup>1</sup>, using some of the characteristic behaviour displayed by OSs when they create packets. A list of OSs identifiable by p0f can be found in Addendum B.

Sudden increases in the amount of traffic received by Windows based OSs would seem to be an excellent indicator of most modern worms as Windows based OSs tend to be the target of such attacks (Moore, Paxson, Savage, Shannon, Stanifor and Weaver, 2003; Moore and Shannon, 2004; NIST, 2008).

Having identified the variables that can be derived from packet captures, the next section discusses how these properties were organized as a database and how analysis was made easier through summarization.

### 3.3 Database Structure and Summarization

The properties mentioned in the previous section were extracted from the packet captures and organized into tables with the primary key being the ID field for a specific packet.

<sup>1</sup><http://lcamtuf.coredump.cx/p0f.shtml>

The packet payload was not included due to space and processing constraints. The organization of these tables can be seen in the entity relationship diagram (ERD) given in Figure 3.1. It was noted that performing simple queries on such a large dataset was time consuming and as such some form of data summarization was required.

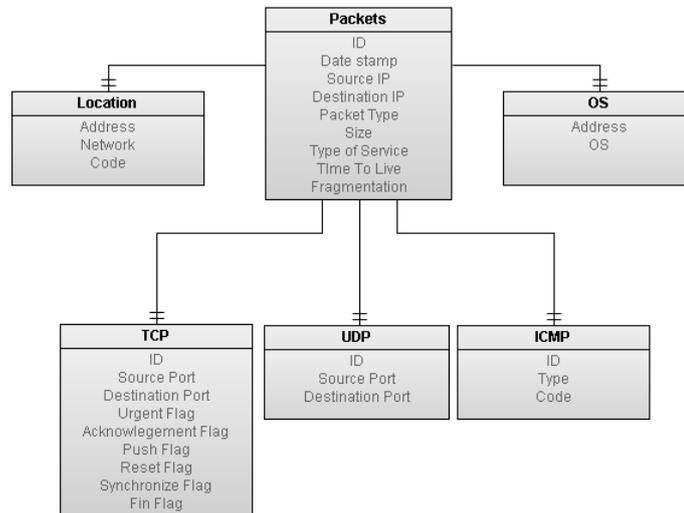


Figure 3.1: ERD of Significant Packet Trace Variables.

Data summarization is an important process in the analysis of network traffic as it reduces large datasets into more manageable components from which more meaningful analysis can be made. averages, medians, deviations and extrema. These statistics included packet counts, cumulative packet size, average packet size, standard deviations in packet size, average TTL, standard deviation in TTL and unique count per hosts at the /32, /16 and /8 level. This data was grouped according to date at the hourly, daily, quarterly and yearly interval and further subdivided according to protocol type with subgroups by port for UDP and TCP and by ICMP type for ICMP. These summarized values were used later by the metrics and techniques in Chapters 5 and 6.

## 3.4 Descriptive Statistics

In this section descriptive statistics, as outlined in Section 2.4.5, are used to understand the nature of the variables significant to this study. It is noted that histograms have been omitted where no value could be gleamed due to intense skewing of counts by dominant factors in the variables.

### 3.4.1 Destination IP

The destination IP variable represents a space which is both large and arguably complex to analysis, though this is dependent on the IP space that the telescope monitors. In this descriptive statistical approach the researchers only considered the overall distribution of packets received per destination IP.

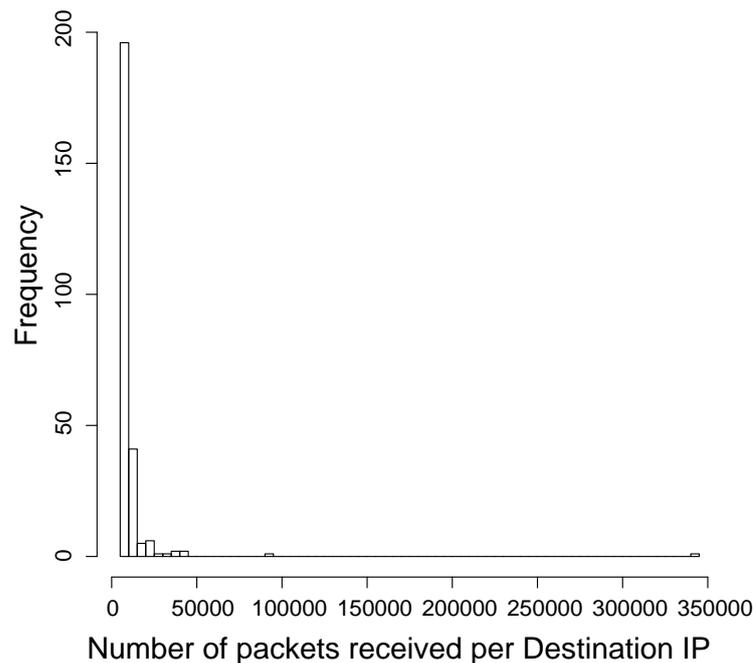


Figure 3.2: Histogram of Destination IP Counts.

The histogram provided in Figure 3.2 clearly shows that the majority of destination IPs received a small portion of the total traffic observed by the network telescope. It logically follows that there must be a number of IPs for which most of the traffic is destined. The observed distribution does seem reasonable as certain IPs would be favoured over others such as in the case of targeted scanning or DDoS backscatter.

Table 3.2 compares the destination IP statistics between the two datasets. Comparatively speaking while the numbers differ there is similar behaviour exhibited by both datasets. Some IPs receive a large contribution of the traffic while a large number of IPs received very little traffic. The deviation in both cases is large which further suggests a large difference between the packets received from a particular destination IP.

| Statistic   | RU1      | RU2      |
|---|----------|----------|
| Ave. Packet Counts /Dest. IP                      | 159382   | 11370    |
| Dev. in Packets Counts /Dest IP                   | 79699.44 | 21985.28 |
| Min Packets Counts /Dest. IP                      | 63695    | 6745     |
| 25 <sup>th</sup> Perc. of Packet Counts /Dest. IP | 90777    | 7415     |
| Median Packet Counts /Dest. IP                    | 147068   | 7834     |
| 75 <sup>th</sup> Perc. of Packet Counts /Dest. IP | 218496   | 9933.25  |
| Max Packets Counts /Dest. IP                      | 731157   | 341600   |

Table 3.2: Descriptive Statistics for Destination IP.

| Space                | IP | Rank      | Count  |
|----------------------|----|-----------|--------|
| IPs $\leq$ a.b.c.127 | 1  | a.b.c.79  | 731157 |
|                      | 2  | a.b.c.88  | 404592 |
|                      | 3  | a.b.c.1   | 396116 |
|                      | 4  | a.b.c.57  | 339544 |
|                      | 5  | a.b.c.37  | 330075 |
| IPs $>$ a.b.c.127    | 1  | a.b.c.252 | 87143  |
|                      | 2  | a.b.c.139 | 87251  |
|                      | 3  | a.b.c.247 | 87327  |
|                      | 4  | a.b.c.134 | 87369  |
|                      | 5  | a.b.c.249 | 87422  |

Table 3.3: Top Destination IPs.

An interesting point of difference between the two datasets is that in the RU1 dataset there is an obvious disparity between counts received by IPs according to which section of the IP range they are found in, as shown in Table 3.3. For RU1 the  $[0,127]$  range received a total of 224719 packets while the  $[128,255]$  range received only 94044 packets, a difference of an order of magnitude. No such disparity exists in the RU2 datasets. This can be explained by earlier versions of Conficker suffering from a flawed host selection algorithm (Irwin, 2011; Porras and Yegneswaran, 2009). Also, it is noted that in both datasets that a.b.c.255 had the lowest counts which is sensible as this IP is reserved for broadcast.

### 3.4.2 Source IP

The source IP variable represents an arguably more complex space to analyse than that of the destination IP space as the source IP is far larger. To deal with this complexity only

the total counts for the source IP together with source IP coverage shall be considered.

Table 3.2 compares summary statistics for the source IP counts for RU1 and RU2. These summaries seem to suggest similar behaviour amongst the dataset with the majority of destination IPs receiving very few packets. This is sensible as the probability of the network telescope receiving a packet from a random IP is low, as was discussed in Section 2.2.1.

| <b>Statistic</b>                                   | <b>RU1</b> | <b>RU2</b> |
|--|------------|------------|
| Ave. Packet Counts /Source IP                      | 7.342      | 9.562      |
| Dev. in Packet Counts /Source IP                   | 275.215    | 168.507    |
| Min Packet Counts /Source IP                       | 1          | 1          |
| 25 <sup>th</sup> Perc. of Packet Counts /Source IP | 2          | 1          |
| Median Packet Counts /Source IP                    | 2          | 1          |
| 75 <sup>th</sup> Perc. of Packet Counts /Source IP | 2          | 1          |
| Max Packet Counts /Source IP                       | 34010      | 46092      |

Table 3.4: Descriptive Statistics for Source IP.

An interesting property of the source IP is the ability to infer the coverage that a particular network telescope has of the Internet. Table 3.5 considers this idea for RU1 and RU2 by listing the source IP coverage for these network telescopes grouped according to different time periods and netblocks. It is clear from the aforementioned table that the coverage of the entire IP space and even considering the /24 space is very small for any given time period. This is to be expected, as the majority of IPs are unlikely to generate malicious traffic. On the other hand looking at the /16 and /8 space produces a larger covered space. This suggests to the researchers that visual techniques that use counts at the /24 and /32 netblock level may struggle to convey this data in a meaningful way as the majority of the /24 and /32 netblocks generate very little traffic observable by a network telescope.

### 3.4.3 TTL

TTL values of all recorded packets appears to be distributed, as noted in Figure 3.3, into three distinct modes as detailed in Table 3.6. Mode A has the largest proportion of Linux OSs which is consistent with the idea that many builds of Linux default the TTL value to 64. The portion of Windows based OSs may be attributed to packets that have travelled a long logical distance. Mode B is dominated by Windows OSs which is consistent with the fact that packets generated from Windows OSs have a TTL of 128.

| Space | Time Unit | Average % Observed     | Maximum % Observed     |
|-------|-----------|------------------------|------------------------|
| /32   | Minute    | $2.132 \times 10^{-7}$ | $7.670 \times 10^{-5}$ |
|       | Hours     | $8.435 \times 10^{-6}$ | $8.022 \times 10^{-5}$ |
|       | Days      | $1.568 \times 10^{-4}$ | $8.810 \times 10^{-4}$ |
|       | Week      | $9.064 \times 10^{-4}$ | $5.124 \times 10^{-3}$ |
|       | Months    | $3.431 \times 10^{-3}$ | 0.019                  |
|       | Year      | $3.025 \times 10^{-2}$ | 0.112                  |
| /24   | Minute    | $7.188 \times 10^{-5}$ | $9.034 \times 10^{-3}$ |
|       | Hours     | $2.804 \times 10^{-3}$ | $1.944 \times 10^{-2}$ |
|       | Days      | 0.047                  | 0.260                  |
|       | Week      | 0.212                  | 1.044                  |
|       | Months    | 0.565                  | 2.295                  |
|       | Year      | 2.592                  | 6.041                  |
| /16   | Minute    | 0.0159                 | 1.565                  |
|       | Hours     | 0.494                  | 2.999                  |
|       | Days      | 3.764                  | 13.281                 |
|       | Week      | 9.0722                 | 22.240                 |
|       | Months    | 15.658                 | 29.257                 |
|       | Year      | 28.651                 | 39.433                 |
| /8    | Minute    | 2.955                  | 36.199                 |
|       | Hours     | 23.031                 | 55.656                 |
|       | Days      | 45.588                 | 72.398                 |
|       | Weeks     | 55.436                 | 77.828                 |
|       | Months    | 60.461                 | 84.615                 |
|       | Year      | 71.041                 | 95.475                 |

Table 3.5: Table of IP by Time Period and IP Class.

Finally, Mode C has a considerable amount ICMP traffic which may be thought to be related to NMAP<sup>2</sup>scanning.

| Mode | Range   | % of Total Data | Related OSs/Application |
|------|---------|-----------------|-------------------------|
| A    | 25-63   | 16.463          | Linux 2.6               |
| B    | 70-127  | 80.418          | Windows XP, 2000 and 98 |
| C    | 230-250 | 2.329           | NMAP Scanning           |

Table 3.6: TTL Modes.

Table 3.7 shows that the distribution of TTL values in the two datasets is extremely similar, through the difference in 25<sup>th</sup> percentile suggests that RU2 received a higher proportion of traffic from Linux based OSs.

---

<sup>2</sup><http://nmap.org/>

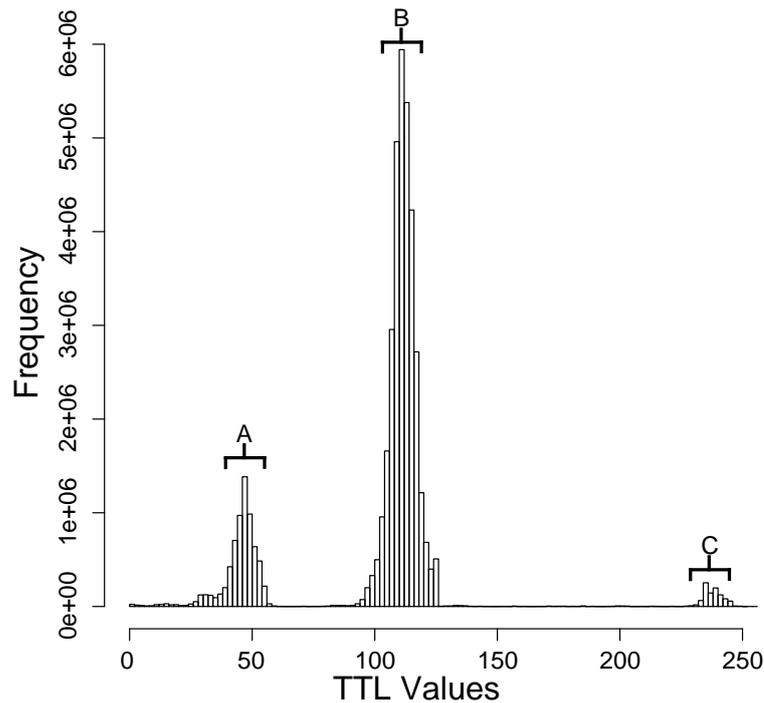


Figure 3.3: Histogram of TTL Counts.

| Statistic                               | RU1    | RU2    |
|---|--------|--------|
| No. of Unique TTL values                | 255    | 211    |
| Ave. Packet Counts                      | 103.9  | 92.862 |
| Dev. in Packet Counts                   | 33.317 | 32.283 |
| Min Packet Counts                       | 1      | 1      |
| 25 <sup>th</sup> Perc. of Packet Counts | 106    | 53     |
| Median Packet Counts                    | 111    | 111    |
| 75 <sup>th</sup> Perc. of Packet Counts | 115    | 115    |
| Max Packet Counts                       | 255    | 253    |

Table 3.7: Descriptive Statistics for TTL.

### 3.4.4 Packet Size

The packet size property is often characteristic of a particular type of behaviour due to the fixed packet size of certain attacks. For example, SQLSlammer infection packets are always 418 bytes in length as the worm makes no attempt to pad its packets.

There appears from Figure 3.4 that there are two modes in the data with some outliers at higher values. Mode A is found between 60-100 bytes, mode B is between 450-500 bytes with outliers found close to 1500 bytes. It could be suggested from this distribution that

the network telescope received a number of attacks that are characterized by low packet sizes such as scanning.

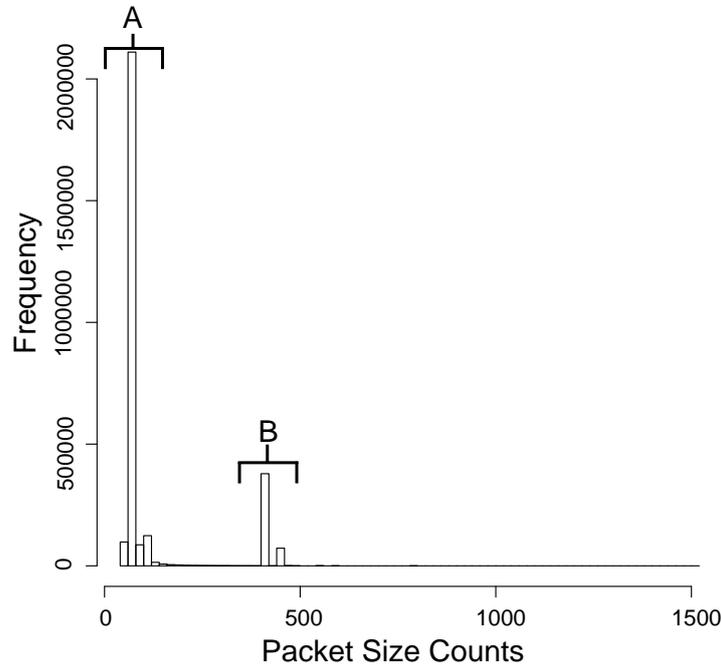


Figure 3.4: Histogram of Packet Size.

A comparison between the packet sizes observed by RU1 and RU2, as provided in Table 3.8, seems to conclude that the two datasets are fairly similar. An interesting note is that the minimum packet size is below 62 bytes which suggests that malformed packets are present in the datasets.

| <b>Statistic</b>                      | <b>RU1</b> | <b>RU2</b> |
|---------------------------------------|------------|------------|
| Ave. Packet Size                      | 101.8      | 128.2      |
| Dev. in Packet Size                   | 120.3749   | 128.320    |
| Min Packet Size                       | 60         | 56         |
| 25 <sup>th</sup> Perc. in Packet Size | 60         | 62         |
| Median Packet Size                    | 111        | 75         |
| 75 <sup>th</sup> Perc. in Packet Size | 156        | 80         |
| Max Packet Size                       | 1514       | 1514       |

Table 3.8: Descriptive Statistics for Packet Size.

Table 3.9 lists the high frequency packet sizes as observed by RU1 and RU2. In both cases a size of 62 had the highest frequency, this is most likely attributed to the fact that both telescopes are passive and TCP handshakes can never be completed and thus only TCP

SYN packets are ever received. A packet size of 418 is found frequently in both datasets and is shown in Section 4.5 to be related to SQLSlammer. Sizes 955 has a high frequency in RU1 but not in RU2 while size 454 bytes was found to be the exact opposite. These packet sizes are shown in Section 4.3 to be related to net send malware.

| Size | RU1      | RU2    |
|------|----------|--------|
| 62   | 24353251 | 827996 |
| 60   | 4339451  | 46461  |
| 418  | 2990266  | 378532 |
| 955  | 30051    | 1929   |
| 454  | 211      | 42668  |

Table 3.9: Counts by Top Packet Sizes.

### 3.4.5 Source Port

The source port could be considered a less significant variable as an intelligent attacker could easily randomize this property. However, this is often not the case due to the laziness of attackers or a attack requiring a specific source port. Table 3.10 shows similar behaviour between the two datasets but with RU1 being of a much larger magnitude. In both datasets at least 96% of the possible 65536 ports are observed, though many of these ports are rarely observed by the network telescopes as evident by the difference between the 25<sup>th</sup> and 75<sup>th</sup> percentiles.

| Statistic  | RU1      | RU2    |
|--|----------|--------|
| Number of Unique TCP Source Ports                        | 65492    | 63165  |
| Ave. Packet Counts /TCP Source Port                      | 508.203  | 21.06  |
| Dev. in Packet Counts /TCP Source Port                   | 9903.123 | 89.918 |
| Min Packet Counts /TCP Source Port                       | 1        | 1      |
| 25 <sup>th</sup> Perc. of Packet Counts /TCP Source Port | 70       | 5      |
| Median Packet Counts /TCP Source Port                    | 91       | 75     |
| 75 <sup>th</sup> Perc. of Packet Counts /TCP Source Port | 118      | 16     |
| Max Packet Counts /TCP Source Port                       | 2334000  | 15440  |

Table 3.10: Descriptive Statistics for TCP Source Port Counts.

Table 3.11 lists the top five TCP source ports for RU1, showing that port 80/tcp appears to be dominant while ports 22/tcp and 6000/tcp make up the majority of packets received. The port 80/tcp dominance may be attributed to the fact that most hosts and routers

are willing to accept packets that have a source port of 80 due to common place nature of HTTP traffic though this does seem unlikely due to most consumer routers having firewalls. It is possible that this 80/tcp is backscatter. Ports 22 and 6000 allow for remote access through SSH (Ylonen, 2006) and X-windows (Scheifler, 1987) respectively so it is sensible that these would have a high proportion of source port counts as remote hackers may attempt to use these applications as a means of entry into a system.

| Port | Count   | % of Total Traffic |
|------|---------|--------------------|
| 80   | 2334289 | 70.134             |
| 6000 | 845593  | 25.405             |
| 22   | 28262   | 0.849              |

Table 3.11: Top TCP Source Ports.

| Statistic  | RU1     | RU2     |
|--|---------|---------|
| Number of UDP Source Ports                               | 65492   | 63165   |
| Ave. Packet Counts /UDP Source Port                      | 302.11  | 15.67   |
| Dev. in Packet Counts /UDP Source Port                   | 378.609 | 128.320 |
| Min Packet Counts /UDP Source Port                       | 1       | 1       |
| 25 <sup>th</sup> Perc. of Packet Counts /UDP Source Port | 2       | 6       |
| Median Packet Counts /UDP Source Port                    | 102     | 4       |
| 75 <sup>th</sup> Perc. of Packet Counts /UDP Source Port | 207     | 116     |
| Max Packet Counts /UDP Source Port                       | 68123   | 15440   |

Table 3.12: Descriptive Statistics for UDP Source Port Counts.

UDP source ports counts shows similar behaviour as TCP ports, high number of ports receive little traffic while a few ports receive a large amount of traffic , but unlike TCP source ports are not as heavily dominated by the top three ports as shown in Table 3.12 and Table 3.13, as even the top five ports make up less than 6% of the total counts. UDP ports 1026/udp, 1027/udp, 1025/udp and 1029/udp are all related to net send malware, a topic explored in Chapter 4.

| Port | Count | % of Total Traffic |
|------|-------|--------------------|
| 1026 | 71136 | 1.381              |
| 1231 | 68071 | 1.322              |
| 1027 | 58546 | 1.137              |
| 1025 | 52385 | 1.017              |
| 1029 | 39339 | 0.764              |

Table 3.13: Top UDP Source Ports.

### 3.4.6 Destination Port

The destination port is a significant variable when identifying malicious attacks as many of the major incidents target weaknesses in specific applications or services that run on home user machines.

The distribution of traffic amongst top TCP ports, as is summarized by Table 3.14, reveals that the majority of TCP destination ports rarely receive traffic. This is inferred from the the high variance and the large difference between the 75<sup>th</sup> percentile and the maximum values in the datasets.

| Statistic   | RU1      | RU2      |
|---|----------|----------|
| No. of Unique TCP Dest. Ports Counts                    | 65168    | 16144    |
| Ave. Packet Counts /TCP Dest. Port                      | 510.729  | 82.38    |
| Dev. in Packet Counts /TCP Dest. Port                   | 67904.13 | 3328.207 |
| Min Packet Counts /TCP Dest. Port                       | 1        | 1        |
| 25 <sup>th</sup> Perc. of Packet Counts /TCP Dest. Port | 5        | 4        |
| Median Packet Counts /TCP Dest. Port                    | 11       | 6        |
| 75 <sup>th</sup> Perc. of Packet Counts /TCP Dest. Port | 1689     | 12       |
| Max Packet Counts /TCP Dest. Port                       | 16893920 | 366625   |

Table 3.14: Descriptive Statistics for TCP Destination Ports.

Table 3.15 list the counts and coverage of the top five TCP ports. From the TCP perspective: ports 445/tcp, 139/tcp, 1433/tcp and 22/tcp are present in all of the months in the RU1 dataset. It would seem sensible to investigate these ports to identify possible incidents.

The distribution of packets amongst UDP destination ports appears to be similar to that of TCP destination ports as suggested by the large variance, high maximum count and narrow inter-quartile range as shown by Table 3.16.

| Port | Packet Counts | % Coverage /hour | % Coverage /day | % Coverage /month |
|------|---------------|------------------|-----------------|-------------------|
| 445  | 16893920      | 39.454           | 98.117          | 99.93 %           |
| 135  | 2749950       | 4.278            | 50.895          | 99.3              |
| 139  | 1680226       | 7.331            | 49.441          | 95.801            |
| 22   | 1224074       | 0.709            | 18.687          | 92.562            |
| 1433 | 156668        | 5.175            | 45.258          | 94.121            |

Table 3.15: Dataset Coverage of TCP Destination Ports.

| Statistic   | RU1      | RU2     |
|---|----------|---------|
| No. of Unique UDP Dest. Ports                           | 65283    | 23532   |
| Ave. Packet Counts /UDP Dest. Port                      | 78.861   | 65.24   |
| Dev. in Packet Counts /UDP Dest. Port                   | 67904.13 | 2722.83 |
| Min Packet Counts /UDP Dest. Port                       | 1        | 1       |
| 25 <sup>th</sup> Perc. of Packet Counts /UDP Dest. Port | 5        | 2       |
| Median Packet Counts /UDP Dest. Port                    | 2070.32  | 6       |
| 75 <sup>th</sup> Perc. of Packet Counts /UDP Dest. Port | 1689     | 13      |
| Max Packet Counts /UDP Dest. Port                       | 2990062  | 378640  |

Table 3.16: Descriptive Statistics for UDP Destination Ports.

| Port | Packet Counts | % Coverage /hour | % Coverage /day | % Coverage /month |
|------|---------------|------------------|-----------------|-------------------|
| 1434 | 2990062       | 77.335           | 99.809          | 100               |
| 137  | 633620        | 0.683            | 25.154          | 95.661            |
| 1026 | 320355        | 0.239            | 11.420          | 50.664            |
| 1027 | 249309        | 0.222            | 10.730          | 50.174            |
| 135  | 22211         | 0.018            | 5.623           | 95.521            |

Table 3.17: Dataset Coverage of UDP Destination Ports.

Table 3.17 list the counts and coverage of the top five UDP ports. Only 1434/udp, 137/udp and 135/udp are observed for the majority of months.

### 3.4.7 ICMP Type

The distribution of ICMP Types as illustrated by the histogram in Figure 3.5 appears to cluster around four distinct modes. The largest mode consists of the majority of the ICMP types with all of the types in this mode having been observed less than 250 times and thus are unlikely to represent a network wide incident or attack. The remaining modes relate to types 0, 3, 8 and 11. Table 3.18 further substantiates this point as both datasets have a large deviation and a large distance between the minimum and maximum values. As has been the case for a number of variables, this behaviour is less pronounced in RU2 which may be attributed to the smaller sampling.

While the two network telescopes see a much wider variety of ICMP types comparative to a normal network node the actual ICMP type coverage is low with TTL times outs, ICMP echoes and Destination Unreachable messages having good coverage over the entire dataset of RU1, as shown by Table 3.19.

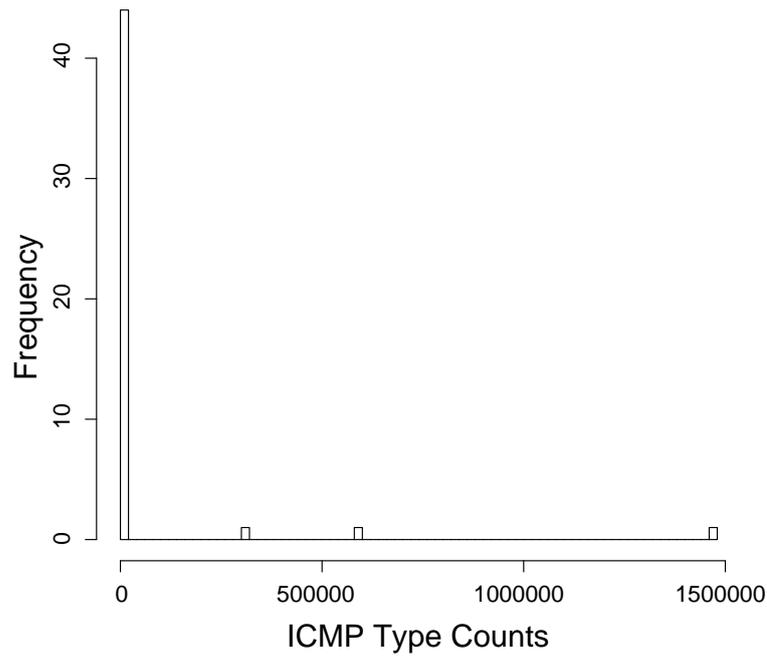


Figure 3.5: Histogram of ICMP Type Counts.

| <b>Statistic</b>                                    | <b>RU1</b> | <b>RU2</b> |
|---|------------|------------|
| Unique ICMP Types Observed                          | 47         | 47         |
| Ave. Packet Counts /ICMP Types                      | 50416      | 5545.62    |
| Dev. in Packet Counts /ICMP Types                   | 32.283     | 232239.2   |
| Min Packet Counts /ICMP Types                       | 1          | 1          |
| 25 <sup>th</sup> Perc. of Packet Counts /ICMP Types | 1          | 19.25      |
| Median Packet Counts /ICMP Types                    | 5546       | 207.5      |
| 75 <sup>th</sup> Perc.of Packet Counts /ICMP Types  | 11321      | 5300       |
| Max Packet Counts /ICMP Types                       | 1471537    | 31202      |

Table 3.18: Descriptive Statistics for ICMP Types.

| Type | % Coverage /hour | % Coverage /day | % Coverage /month |
|------|------------------|-----------------|-------------------|
| 8    | 17.649           | 99.681          | 99.930            |
| 3    | 18.279           | 80.018          | 96.29 904         |
| 11   | 5.132            | 63.820          | 99.37 91          |
| 0    | 0.084            | 1.946           | 22.11 6           |
| 5    | 0.013            | 0.717           | 13.016            |

Table 3.19: Dataset Coverage of ICMP Types.

### 3.4.8 OS Type

Windows OSs have had the majority of the OS market for a lengthy period of time even though this has diminished in recent years (Gruener, 2012; W3Schools, 2012). If this detail is coupled together with the fact that Windows based OSs have had a long history of exploitation by hackers it is expected that the majority of traffic received from a network telescope will be from Windows OSs. The OS identification tool p0f was used to perform OS identification and as such the study will be limited to the accuracy and diversity of p0fs identification.

The proportion of traffic received from specific OS types is provided in Table 3.20. It is clear from this table that during the time period observed by RU1 that Windows XP/2000 based OSs are dominant. While traffic from Linux and BSD hosts was observed by the network telescope, this traffic made up a small fraction of the total data.

| Port                | % of Total Traffic |
|---------------------|--------------------|
| Windows 2000, XP    | 96.411             |
| Windows 98          | 2.114              |
| Windows 2003        | 0.465              |
| Linux 2.4-2.6       | 0.469              |
| FreeBSD 6.x         | 0.092              |
| Linux 2.6           | 0.135              |
| Windows NT 4.0      | 0.032              |
| MacOS 9.1           | 0.027              |
| NetBSD 1.6W-current | 0.026              |

Table 3.20: Top OS Types.

### 3.4.9 Variable Frequency

It may be valuable instead of considering the counts of a particular quantity to rather consider the number of contiguous packets observed by the network telescope. Contiguous packets are thought to share at least one similar property with a packet within the last minute of the following: source IP, destination IP, source port, destination port and packet size. A simple justification for this would follow from considering the number of contiguous packets in the dataset as shown in Table 3.21, which shows the percentage of packets which share a particular property with at least two other packets in a five minute window. A number of packets share at least one common property with the exception of

source port and source IP. From this one may conclude that the traffic observed is really a series of attacks with a sampling of ‘random’ packets interspersed between attacks.

| Shared property        | % of packets |
|------------------------|--------------|
| Source IP              | 17.987       |
| Destination IP         | 67.69        |
| Packet Size            | 94.205       |
| Source Port (TCP)      | 82.114       |
| Source Port (UDP)      | 27.539       |
| ICMP Type              | 78.122       |
| Destination Port (TCP) | 92.522       |
| Destination Port (UDP) | 78.807       |

Table 3.21: Packets with Related Characteristics.

## 3.5 Summary

This chapter has provided background and statistical insight on the network telescope datasets used in this thesis. Section 3.1 documents the origin of the two major datasets RU1 and RU2 as being sourced individually from separate /24 network telescopes. The location, period covered and number of packets captured for each datasets was also provided.

The significant variables derived from the datasets were discussed in Section 3.2. A brief description of the relevant properties of each of the variables was provided together with the significance that, that particular variable may hold when attempting to identify malicious activity in network telescope datasets. The variables described included: source and destination IP, packet size, datagram type, ICMP type or TCP/UDP port, TTL and OS type of the host machine.

The summarization of the dataset was discussed in Section 3.3 as a way to deal with the large size of the datasets. These summarization fields greatly reduce the time taken to analyse portions of the datasets and were used by many of the calculation in future chapters.

A short descriptive statistical study is made of the significant network telescope variables in Section 3.4. A number of interesting properties important to later work were discovered such as the low coverage of the IPv4 space observed by the network telescopes, the tri-

modal nature of TTL, the distribution of source and destination ports and the majority of traffic sampled originating from Windows hosts.

Having studied the datasets in a broad sense it is now necessary to examine the datasets for major incidents that may be used as test cases for analysis techniques used in later chapters. The following chapter listing a number major incidents identified in the RU1 dataset and their characteristic behaviour.

*Attacks whether criminal or not, are exceptions. They're events that take people by surprise, that are 'news' in its real definition.*

Bruce Schneier, *Secrets and Lies*

# 4

## Notable Incidents

THE concept of incident in this thesis refers to an exceptional case found in the data which sheds light on potential outbreaks or indicates a change in the composition of malicious traffic on the Internet. As the majority of traffic received is in some form malicious, attempting to identify each case would require a considerable amount of human-based analysis and research. While a number of well documented cases of worm outbreak and other malicious activity have been recorded by reputable organizations, such as the US Computer Emergency Response Team (CERT)<sup>1</sup>; SANS (System Administration, Networking and Security Institute)<sup>2</sup> and independent information security researchers, it is important to consider what incidents are observable based solely upon the characteristic behaviour found in the data collected by a network telescope.

The limitations of observable activity by network telescopes has already been discussed in Section 2.2 and may be summarised by stating that network telescopes can only observe incidents that are directed towards the network telescope specifically, towards the Internet as a whole or when the network telescopes IP space has been spoofed. Further, as the

---

<sup>1</sup><http://www.us-cert.gov>

<sup>2</sup><http://www.sans.org/>

network telescopes used were passive some of the specific characteristics of a particular incident may not be observable. It is thus accepted that the incident classification provided is only partial and could potentially be inaccurate due to the passive nature of the network telescopes used, the complex nature of the datasets, the interplay of other incidents and other external factors which complicate the process

A description of a selection of the significant incidents identified in the 50 months of data obtained from the RU1 dataset is provided in this chapter. The remainder of this chapter is structured as follows:

- Section 4.1 briefly describes how network telescopes can observe DDoS and identifies five cases of DDoS observed via backscatter.
- Section 4.2 discusses the properties of Conficker and shows evidence to suggest that the five Conficker variants are present in the network telescope dataset.
- Section 4.3 investigates malicious net send messages that were observed in the network telescope data.
- Section 4.4 defines the types of scanning activity often observed on the Internet and considers cases found in the network telescope datasets.
- Section 4.5 discusses how some incidents are still active long after initial infection by considering the SQLSlammer related traffic still visible in the RU1 dataset.
- Section 4.6 concludes this chapter by summarizing the major points discussed in this chapter.

## 4.1 DDoS Backscatter

The objective of DDoS is to prevent access to a particular service for economic, political or malevolent reasons (Bidgoli, 2006). In the past a number of major organizations have been unable to provide services to their consumer base due to DDoS resulting in a loss of revenue. The DDoS against Sony's PSN network by 'Anonymous' is an example of such an occurrence (Mick, 2009). A common technique employed in DDoS is to spoof an IP range or IPs from a random IP space to protect the perpetrator's identity and then use these IPs to flood a host with requests. If this sudden influx of traffic exceeds the limitations of the victims computing and bandwidth resources, legitimate users experience a lack

of access or degradation in the quality of the service. Once the targeted hosts are able to process and respond to all of the requests made, those requests that were generated by the spoofed IPs can longer be routed to those spoofed spaces. The TTL of these packets expire causing the generation of ICMP type 11 packets which are sent back to the address space that was spoofed. Occasionally it occurs that this address space belongs to a network telescope (Moore, Voelker and Savage, 2006; US-CERT, 2009b). This type of indirect observation is commonly referred to as backscatter (Pang, Yegneswaran, Barford, Paxson and Peterson, 2007; Pemberton, Komisarczuk and Welch, 2007; Weaver, 2010b). It is worth noting that one of the alternative approaches to DDoS where no spoofing occurs, such as cases of DDoS based ‘hactivism’ (Wagenseil, 2011) where individual users make use of applications such as the Low Orbit Ion Cannon<sup>1</sup>, is generally unobservable by a network telescope as no traffic is directed towards the telescope.

The major observable cases of suspected DDoS in RU1 dataset are listed in Table 4.1. Each of the identified cases represent a SYN flood DDoS attack on a web server. This was inferred from the fact that all of the ICMP backscatter from these victim hosts returned the original packet sent from the attackers. These packets all had the SYN flag set and a destination port of 80/tcp. SYN flooding attempts to force a web server to open a number of TCP connections that are only half complete by sending a single packet with a SYN flag, when the server responds with a SYN-ACK the attacker never responds causing this connection to remain half-open (Eddy, 2007). If the attacker can force open enough of these connections legitimate users are unable to form TCP sessions with the server due to the limitations on the number of simultaneous connections that the server OS allows. It is noted that almost all of the major incidents of DDoS appear to be directed towards hosts based in China.

The packet traces for these periods of major DDoS were then investigated by considering the original packet captures observing the exact start and end time and noting any interesting properties. In each of the incidents packets were received from a single source IP this is sensible as victims of the DDoS attack are likely to send all their traffic through a single external IP.

Figure 4.1 depicts the number of ICMP type 11 packets received during DDoS.D. It is clear that before and after the incident there is little to no ICMP type 11 traffic then suddenly a brief period of high volume traffic. It is suspected that metrics that focus on the packet sizes, distribution of source and destination IPs and counts related to ICMP

---

<sup>1</sup><http://sourceforge.net/projects/loic/>

| Incident Name | Start Time          | Suspected Target | No.of packets |
|---------------|---------------------|------------------|---------------|
| DDoS.A        | 2008-02-22 06:02:00 | China Unicom     | 1733          |
| DDoS.B        | 2008-05-26 14:02:17 | China Unicom     | 1674          |
| DDoS.C        | 2008-10-18 17:59:00 | CHINANET         | 1534          |
| DDoS.D        | 2009-02-17 21:30:49 | China Unicom     | 159990        |
| DDoS.E        | 2009-05-26 21:30:49 | China Unicom     | 2490          |

Table 4.1: Major DDoS Incidents Observed.

type 11 would be effective in identifying this kind of behaviour.

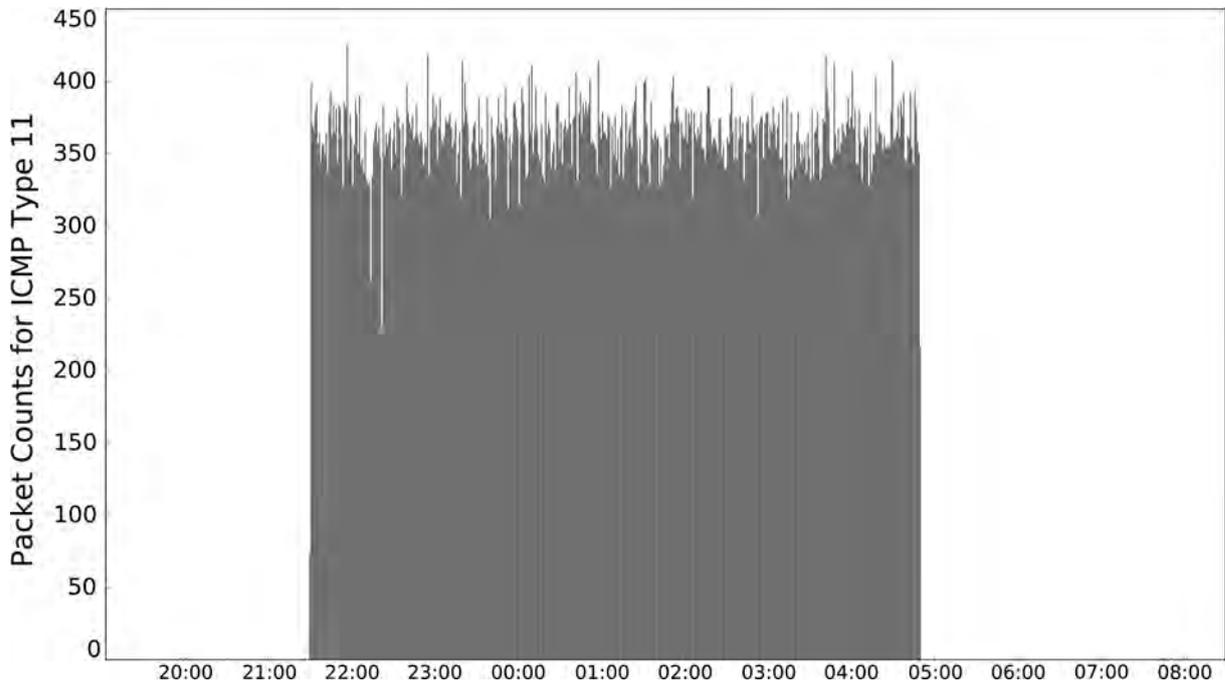


Figure 4.1: Counts for ICMP type 11 during DDoS.D (/minute).

## 4.2 Conficker

The Conficker worm, also known as DownAdUp, appeared in late 2008 and has become one of the most prevalent worms to infect machines across the Internet due to its technical versatility (Nahorney, 2009). Conficker exploited vulnerabilities in the Microsoft Windows Remote Procedure Call (RPC) stack through specially crafted RPCs over port 445/tcp (Microsoft, 2008; Porras and Yegneswaran, 2009). Conficker initially made use of a number of secondary “tricks” such as brute forcing network passwords, infection

through removable media such as USB flash drives and encrypted payloads (Porras, Saidi and Yegneswaran, 2011). These techniques ensured the worm spread far and protected itself from external takeover (Nahorney, 2009). It is theorized that the primary goal of the worm was to construct a large botnet that could be sold to perform illegal activities (Weaver, 2010a). As Conficker exploits a vulnerability in the RPC stack, a host needs to be able to complete a three-way handshake in order to receive the payloads. This is a significant caveat for analysis from passive network telescopes.

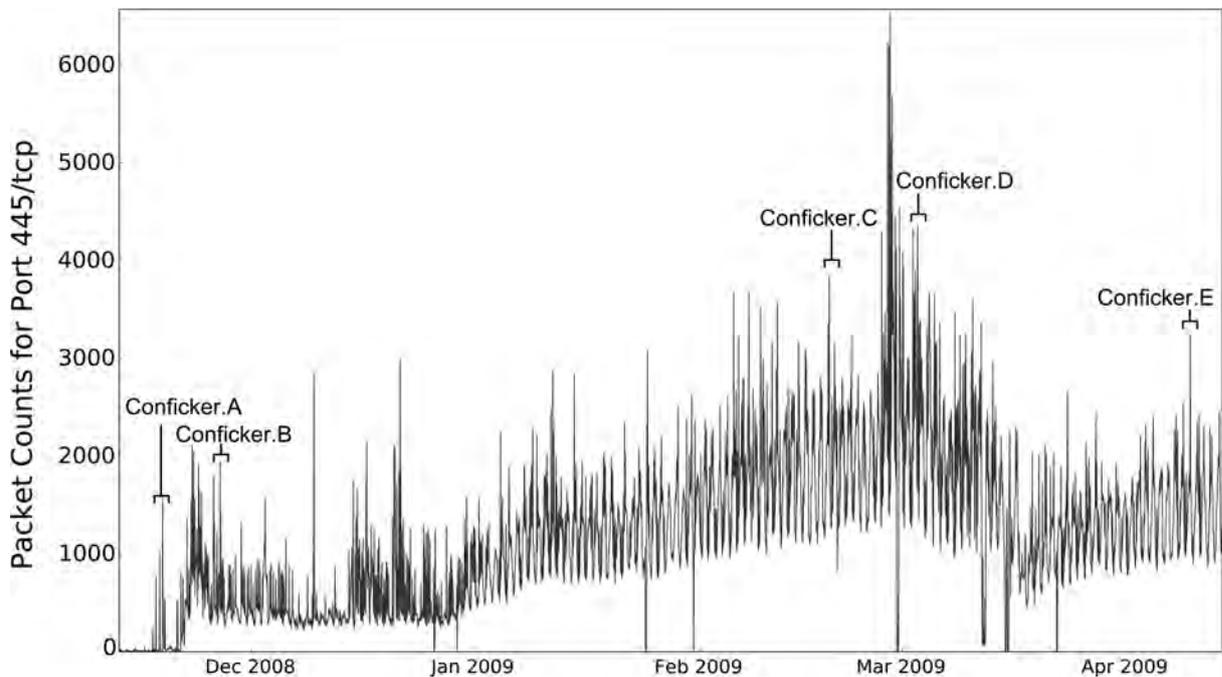


Figure 4.2: Packet Counts for Dest. Port 445/tcp (/hour).

The Conficker Working Group<sup>1</sup> has created a comprehensive time-line of the important events in the growth of the Conficker worm as agreed upon by a collection of experts involved in the analysis of Conficker (Conficker Working Group, 2009). Summarizing this time-line in terms of the emergence of the major variants of Conficker yields Table 4.2. Figure 4.2 shows how these variants were observed by the network telescope together with the approximate outbreak of each of the variants listed in Table 4.2. Conficker.A occurs at a time where the volume of 445/tcp traffic previously received was relatively stable. Conficker.A itself represents a significant spike in the amount of 445/tcp traffic received by the telescope. Conficker.B represents another fairly significant spike in 445/tcp following after a brief period of lower activity trailing a large spike in 445/tcp caused possibly by further scanning or hosts infected by Conficker.A attempting to infect the the IP space of the network telescope. Conficker.C marks a minor local extrema in the dataset.

<sup>1</sup><http://www.confickerworkinggroup.org/wiki/>

| Variant Name | Date             |
|--------------|------------------|
| Conficker.A  | 20 November 2008 |
| Conficker.B  | 28 December 2008 |
| Conficker.C  | 20 February 2009 |
| Conficker.D  | 4 March 2009     |
| Conficker.E  | 8 April 2009     |

Table 4.2: Emergence of Conficker Variants.

Conficker.D also marks a minor extrema in the dataset following on from a major extrema in the dataset. Conficker.E occurs after some stabilization of the received 445/tcp traffic and is represented by a fairly large spike in the data.

### 4.3 Windows Net Send Malware

Socially engineering users into following links and installing malware on their systems is an effective delivery mechanism for malware. A popular technique in early 2004 and 2005 was to make use of the the Windows Messenger Service through net send messages to deliver these malware links (Microsoft, 2001). This was effective against many home users who in general lack knowledge in information security. These messages attempt impress that their PCs are currently insecure or infect with a virus. The message will then usually provide a URL that the user is encouraged to follow to remediate the issue. This usually results in the downloading of malicious files or the suggestion to purchase faux anti-virus software (Microsoft, 2007). Figure 4.3 shows a sample of such a net send message. These net sends messages tend to be sent by a single host to every IP in the telescope IP in a non-sequential order to both 1027/udp and 1026/udp. For a particular attack the same message is typically used implying the same packet size. This packet size tends to be very large as these messages are usually long, as Figure 4.3 exemplifies. A total of 135 unique messages types were observed in the RU1 network telescope dataset.

Figure 4.4 plots packet counts observed for destination ports 1026-1027/udp the entirety of RU1. It is worth noting that this exploit only affect Windows XP/2000/2003 machines as Microsoft Vista onwards no longer support this service. Net send malware is prevalent during the 2005 to 2007 period with a diversity of messages. Almost no traffic is observed for destination ports 1026-1027/udp after this point until net send malware reappears in August 2008. Between August 2008 and June 2009 only three new message types are

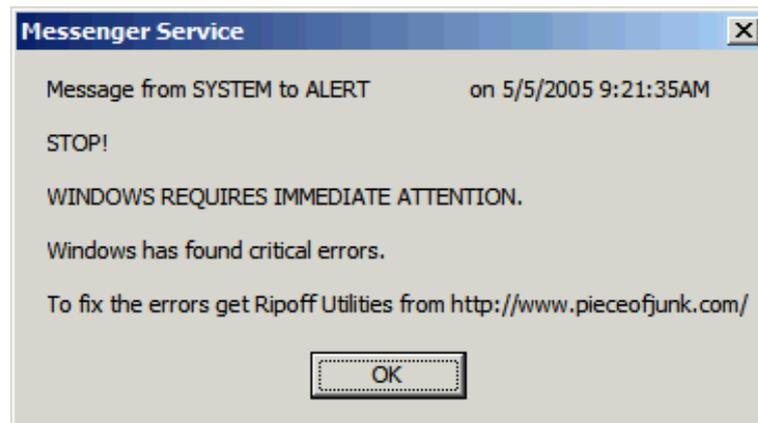


Figure 4.3: A Hoax Net Send Message.

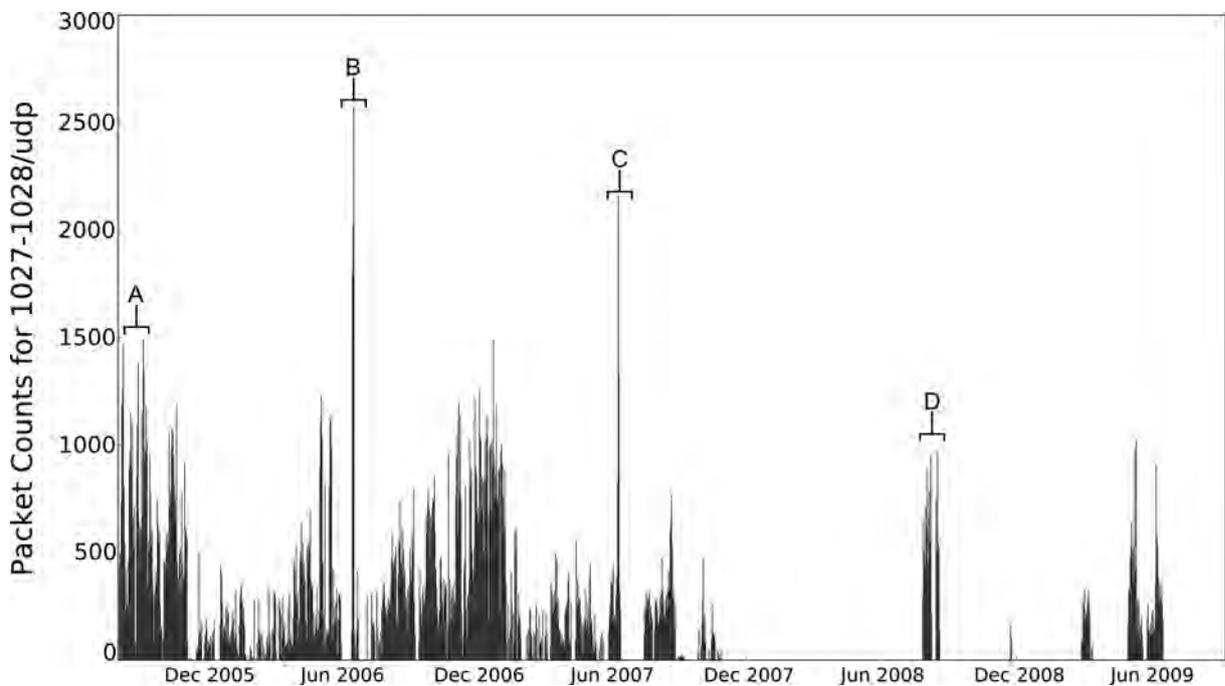


Figure 4.4: Packet Counts for Dest. Ports 1026-1027/udp (/day).

observed after which net send malware appears to have vanished. Table 4.3 lists four of the cases of net send malware considered in this thesis, these points are labelled on Figure 4.4.

| Variant Name  | Packet Size | Counts |
|---------------|-------------|--------|
| MessMalware.A | 991         | 18149  |
| MessMalware.B | 557         | 20566  |
| MessMalware.C | 611         | 4396   |
| MessMalware.D | 636         | 34923  |

Table 4.3: Cases of Net Send Malware.

## 4.4 Scanning Activity

Port scanning segments of the Internet is a popular reconnaissance technique amongst hackers and ‘script kiddies’ to find vulnerable hosts on the Internet which may be used to seed worms or be exploited in other means. There are number of forms of scanning that vary the depth, breadth and period over which the scanning occurs. In this thesis of primary interest shall be the depth and breadth approaches as long time period scanning can be difficult to detect (Van Riel and Irwin, 2006a; Van Riel, 2006b). Depth scanning is defined in this thesis as scanning that targets a specific port across a number of hosts. Two prominent cases of this behaviour found in RU1 are 137/udp and 22/tcp scanning. Port 137/udp is associated with NetBIOS which showed intensified scanning in 2000 that has continued on to the present day (CERT, 2000; Wireshark Forums, 2011). Port 22/tcp scanning is sensible in the modern networking context due to ubiquity of Secure Shell (SSH) which hackers target as a means to gain access to a particular host. Table 4.4 details the properties of a case of 137/tcp and a case of 22/tcp scanning, these incidents are used to evaluated the techniques studied in this thesis in Chapter 7.

| Variant Name | Date             | Port     | Counts |
|--------------|------------------|----------|--------|
| Scan.A       | 13 December 2006 | 6000/tcp | 371    |
| Scan.B       | 14 Jan 2008      | 137/udp  | 9626   |

Table 4.4: Cases of Depth Scanning.

Breadth based scanning attempts to gain as much information from a particular network by scanning a range of ports to find potential weaknesses. Two cases of this activity observed in the RU1 dataset are listed in Table 4.5.

| Variant Name | Date            | Port Range Scanned | Packet Count |
|--------------|-----------------|--------------------|--------------|
| Scan.C       | 15-19 June 2006 | 0-6000             | 9727         |
| Scan.D       | 2-4 Jan 2007    | 7000-25000         | 20566        |

Table 4.5: Cases of Breadth Scanning.

## 4.5 SQLSlammer

It was mentioned in the introductory chapter that one of the issues with analysing network telescope data is the persistence of old exploits such as SQLSlammer. While SQLSlammer data is not used as a case study of an incident, it was filtered out from the raw data for the formation of metrics as it only served to hinder analysis. The vector used by SQLSlammer could be identified as destination port 1434/udp and a specific 418 byte payload. The specifics of this payload and infection mechanism are further explored in (Ray, 2004). Figure 4.5 illustrates that SQLSlammer is still a component of the traffic observed by network telescopes, contributing a number of the spikes that may be observed by metrics derived from this data, though generally speaking SQLSlammer traffic is a very small contributor.

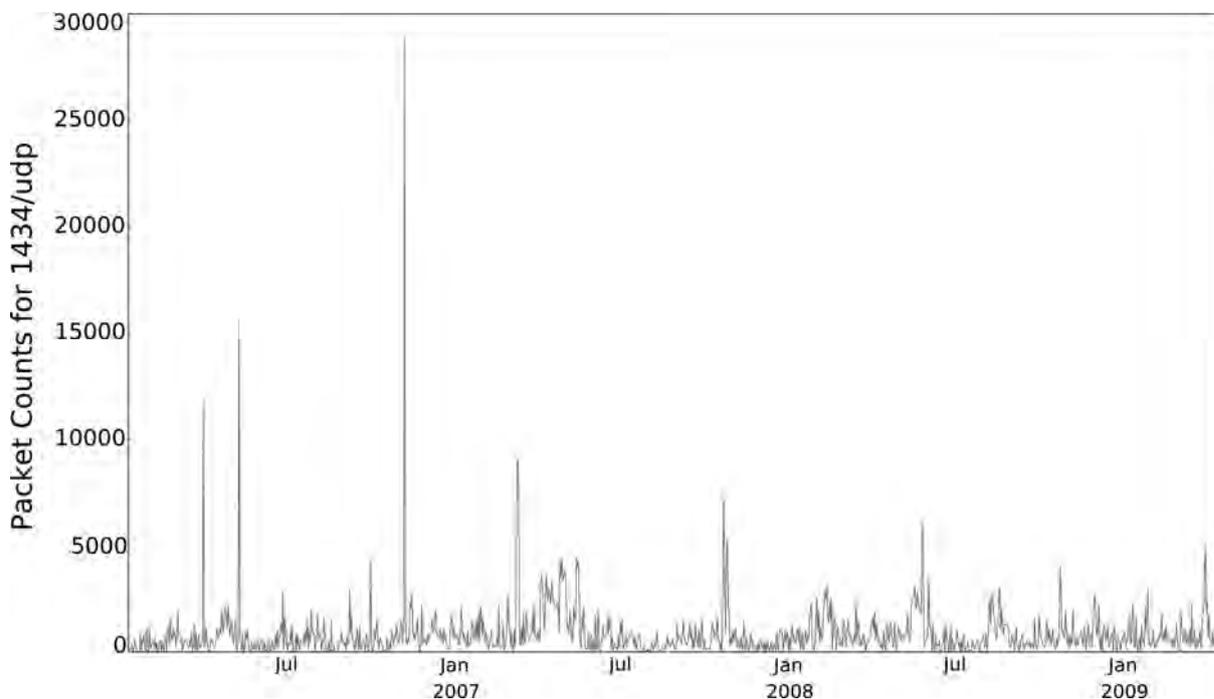


Figure 4.5: Packet Counts for Dest. Port 1434/udp (/day).

While SQLSlammer is unlikely to become a major threat to the functioning of the Internet, as a small percentage of hosts are still vulnerable to these exploits, it is worth mentioning for the effect that this traffic has on metrics derived from network telescope datasets.

## 4.6 Summary

This chapter has identified and described the characteristics of some of the major points of interest observable in the RU1 dataset.

DDoS is a tactic which overloads a server with more traffic than it can possibly process, preventing legitimate users from accessing the services provided by said server. DDoS observed via backscatter was discussed in Section 4.1 by describing the notion of backscatter and listing five cases observed in the RU1 dataset.

The outbreak of the Conficker was contemplated in Section 4.2 together with the effect of observing this behaviour from a passive network telescope perspective. The outbreak of five of the major DDoS variants is discussed with respects to changes observed in the packet counts for destination port 445/tcp.

Another interesting incident noted in the network telescope dataset was the presence of malicious net send messages as discussed in Section 4.3. These net send messages attempt to “trick” a naive user into following links to malicious files or attempt to convince users to buy fraudulent anti-malware applications. These messages were characterized by spikes in traffic destined for ports 1026/udp and 1027/udp and bursts of packets with the same size.

Port scanning is prolific on the Internet and makes up an unsurprisingly large component of the network traffic observed by the network telescope. Section 4.4 groups scanning into breadth or depth scanning types and considers some ports commonly associated with scanning.

SQLSlammer, a worm that is commonly thought to pose no significant threat to the state of the Internet, is still observed by network telescopes seven years after initial infection. Section 4.5 briefly discusses the properties of SQLSlammer and how it affects the analysis of network telescope data.

Having identified a number of incidents present in RU1 it is now possible to consider

potential metrics and analysis approaches. The following chapter considers the formation and behaviour of metrics created through counts, aggregation and groupings of network telescope variables.

*If you cannot measure it you cannot control it.*

Tom DeMarco, Software Engineer paraphrasing Lord Kelvin

# 5

## Metrics for Incident Identification

**M**ETRICS provide insight into the state of a particular entity by the comparison of current measurements of the observable components of said entity against previously observed measures. Metrics have been used in a number of monitoring contexts such as employee performance, service delivery and equipment maintenance. Within the information security context metrics have been used to measure the effectiveness of information protection strategies and to estimate the risk that an organization faces with regards to data loss or theft (Jaquith, 2007).

As metrics provide a context for understanding the change in an entity it seems sensible that metrics derived from network telescope data may be of value in identifying new trends in the malicious activity observed on the Internet. The proportion of TCP datagrams to total traffic, the number of packets received with a specific destination port and the average packet size are examples of metrics that may be of interest to analysts when using network telescope datasets to identify new trends in malicious activity.

This chapter explores metrics and related issues within the network telescope domain according to the following structure:

- Section 5.1 considers the normality of data obtained from a network telescope and the difficulties involved with such a definition. An example of how normality is useful in network traffic analysis is illustrated by studying the effect that Conficker had on the total packets counts received per month by a network telescope.
- Section 5.2 explores metrics which can be formed through simple counts of quantities. In particular total packet counts, counts per specific packet size, counts per specific destination IP and counts per port or ICMP type are considered.
- Section 5.3 aggregates a selection of the network telescope variables through averages and extrema. The following metrics formed from this methodology were studied as examples: average packet size, average TTL value, average counts per source IP, average counts per destination IP and maximum counts per destination IP.
- Section 5.4 investigates a grouping approach of the significant variables according to specific criteria. The metrics investigated include: counts per protocol, protocol ratio's, counts by OS type and bogon IP counts.
- Section 5.5 concludes the findings in this chapter and summarizes the results found.

A portion of the work considered in this chapter is based on a paper written for the Southern African Telecommunication and Applications Conference (SATNAC) Conference in 2010 (Cowie, 2010b).

## 5.1 Normality for Network Telescope Metrics

In order for metrics to be effective, there is a need to be able to identify unexpected behaviour. To be able to achieve this a definition of normality is required, though this is a challenging task. The major issues with such a definition include: the vast quantity of traffic to be analysed and classified, dilution of incidents due to interplay with other incidents and frequent changes in both the physical and logical composition of the Internet. These factors make it difficult to justify a given expectation as there is no set of rules that could adequately explain the nature of the malicious activity on the Internet as this is subject to human influence, undisclosed vulnerabilities and other factors which could not be known ahead of time.

One approach is to consider the normality of network traffic to be composed of a number of measures associated with network activity which derive an expected value based upon

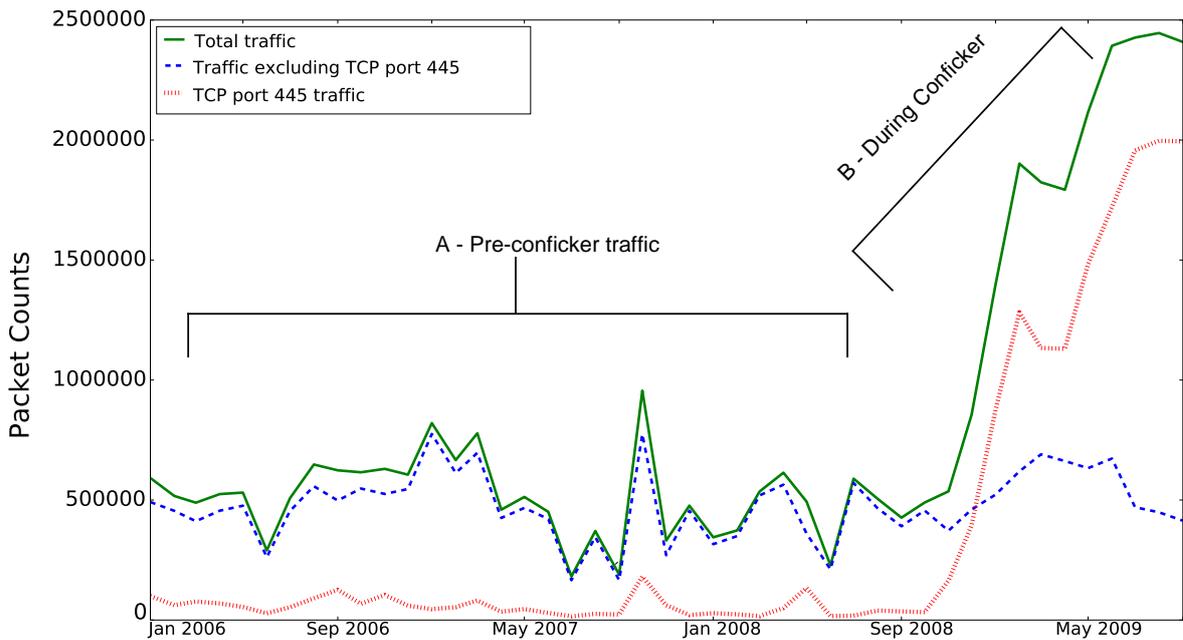


Figure 5.1: Total Packet Counts (/month)

historical values. Packet counts per port, ratio of packets per port and the distribution of packet sizes are a few examples of such norms. An illustrative example of the violation of normality can be seen in Figure 5.1. It is clear from traffic marked at A on the figure, pre-Conficker (21 October 2008), that there is a definite relation between the packet counts of traffic inclusive and exclusive of destination port 445/tcp. That is to say that the packet counts peak and dip at relatively the same place and are generally similar in shape. This implies that while 445/tcp traffic is a component of the total traffic, it is not the major contributor of the traffic received. This creates a precedent for the composition of malicious traffic to not be dominated by 445/tcp traffic, based upon data collected for the last 75 months. From the section marked as B on the figure onwards there is logistic growth in both the total traffic observed and 445/tcp traffic while the traffic exclusive of port 445/tcp continues in a similar fashion as previously observed. The total traffic now appears to be dominated by the behaviour of 445/tcp and this suggests that something unusual has occurred.

## 5.2 Analysis by Counts

Counts of the major variables observed in the datasets provides a simple and fast measure of normality by comparison with previously measured counts for similar periods. This section describes the following metrics that can be formed purely from counts: total packet counts, counts per specific packet size, counts per specific destination IP and counts per specific port or ICMP type.

### 5.2.1 Total Packet Count

The total packet count metric is simple in its definition: it is the count of every packet received by the network telescope for a particular time grouping, such as per hour or per day. Figure 5.2 considers the total packet count at a per day level together with a selection of the notable spikes and dips marked. The total packet count metric seems to act as an effective ‘high level’ metric as it is capable of observing the larger trends such as Conficker in the dataset. However, when it comes to the lower volume incidents it is difficult to discern their impact on the overall packet count. This issue is likely to be compounded by larger network telescopes due to the far larger quantity of traffic observed.

The total packet count metric appears to hover around approximately 20000 packets per day but undergoes significant change during the outbreak of Conficker, showing logistic type growth commonly associated with viral outbreaks (Qiming, Xu and Wang, 1999; Weaver, 2010a). Table 5.1 highlights some of the major points in Figure 5.2 together with the expected causes of this activity. It appears that the total packet count metric can identify a variety of incidents from net send malware to increased scanning activity as long as the incident is sufficiently large. A couple of cases of depth port scanning, breadth scanning and net send malware can also be observed from this metric.

| <b>Incidents</b> | <b>Cause</b>                          |
|------------------|---------------------------------------|
| A                | Net Send Malware                      |
| B                | TCP Port Scanning Activity            |
| C                | Breadth scanning                      |
| D                | Increase in activity on port 1038/udp |
| E                | Net Send Malware                      |
| F                | Conficker                             |

Table 5.1: Incident Listing for Figure 5.2 (Total Packet Count).

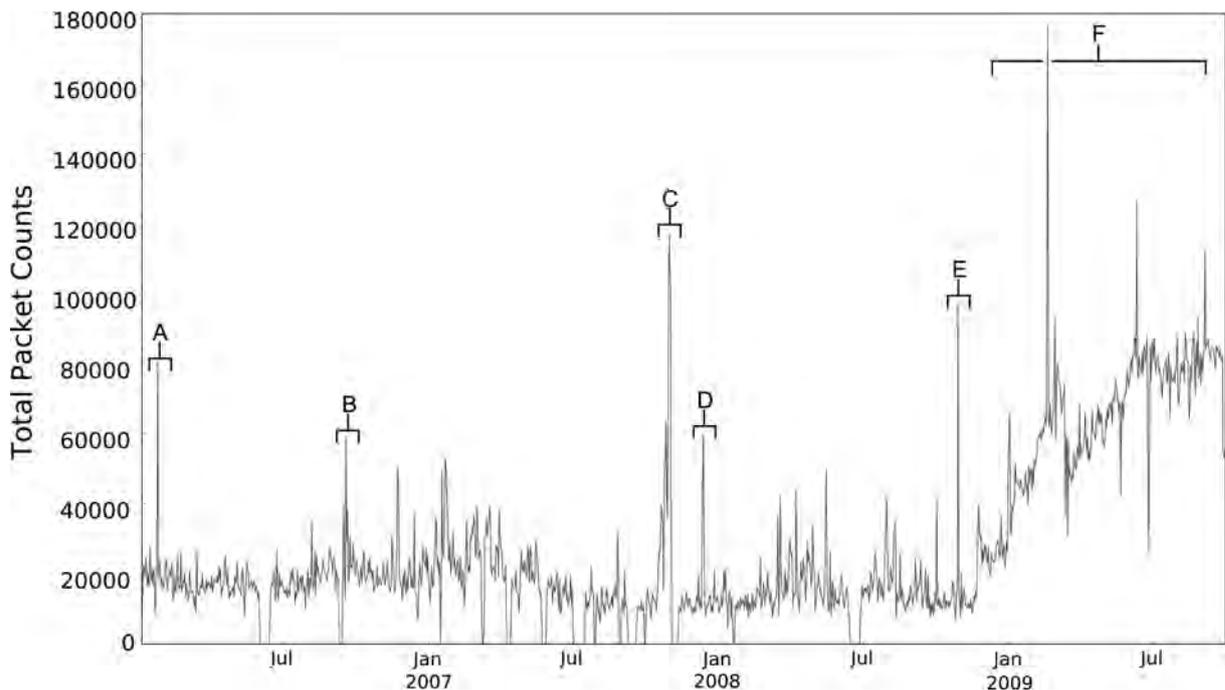


Figure 5.2: Total Packet Counts (/day)

### 5.2.2 Counts by Specific Packet Sizes

Counting the frequency of particular packet sizes is an unusual metric but yields some interesting results. In the RU1 dataset there are 809 distinct packet sizes, which is reasonable if it is considered that any packet size is possible with the range [62,1500] bytes. The issue becomes how to select which of those 809 frequencies to display as it not possible to display all of those frequencies at once due to data overload and further the graph would be difficult to read. One realization that aids in this selection is that some packet sizes are more significant than others. Based on this hypothesis Figure 5.3 was created, with the significant packet sizes chosen as 60, 62, 418 and 955 bytes, which were shown to be of interest in Section 3.4.4.

Packets of 60 bytes or less are malformed as they do not meet the minimum requirements for a IPv4 packet as outlined in RFC 791 (Information Sciences Institute, 1981a), however it is possible to generate packets of this size this through the use of NMAP and other DDoS tools. Point C was identified as a case of major DDoS and substantiates this argument. Packets that are 62 bytes long are interesting as they represent the smallest possible legitimate packet that can be sent and as a result packets received by a passive network telescope are most likely to be 62 bytes due as TCP packets will be generally of this length.

Point D illustrates this point as it shows the characteristic logistic growth associated with Conficker. The size 418 bytes is another significant packet size as it is the size of SQLSlammer infection packets, which was discussed in Chapter 4, being a contributor to network telescope datasets even though SQLSlammer is no longer considered a threat to the Internet. Finally, 955 length packets show long periods of activity followed by periods of no recorded activity which was shown to be related to net send malware in Section 4.3.

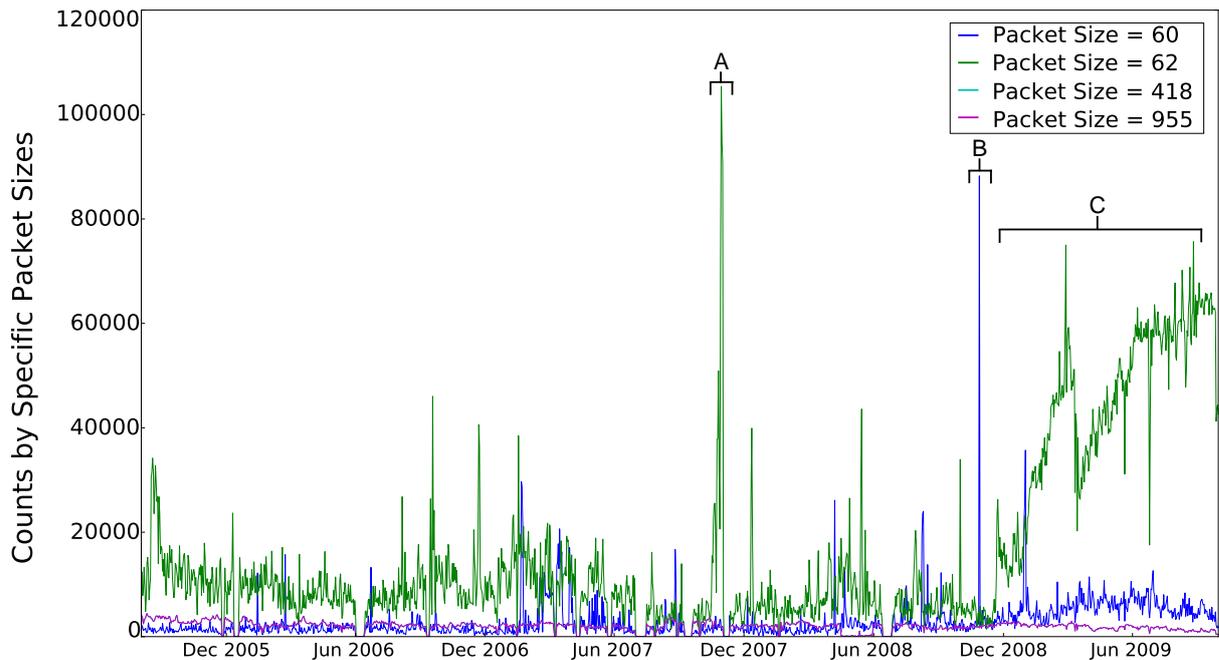


Figure 5.3: Packets Counts per Packet Size (/day).

### 5.2.3 Destination IP Counts

Destination IP counts suffers from the issue, like packet counts and TTL counts, that there are simply too many destination IPs to consider and display each individually. Picking destination IPs based on significance is more challenging than is the case for packet sizes as there is little significance outside of IPs reserved for broadcast (a.b.c.255) and the network identifier (a.b.c.0) though a case could be made for IPs commonly associated with functions such as the default gateway (a.b.c.1). To avoid this issue a selection of four destination IPs were chosen to study the influence of Conficker's flawed propagation algorithm as an example application of this metric. The chosen IPs were a.b.c.43, a.b.c.67, a.b.c.197 and a.b.c.207 counted at a per day level as shown in Figure 5.4.

It is clear from the points A, B and C in Figure 5.4 that the spikes and dips in the counts

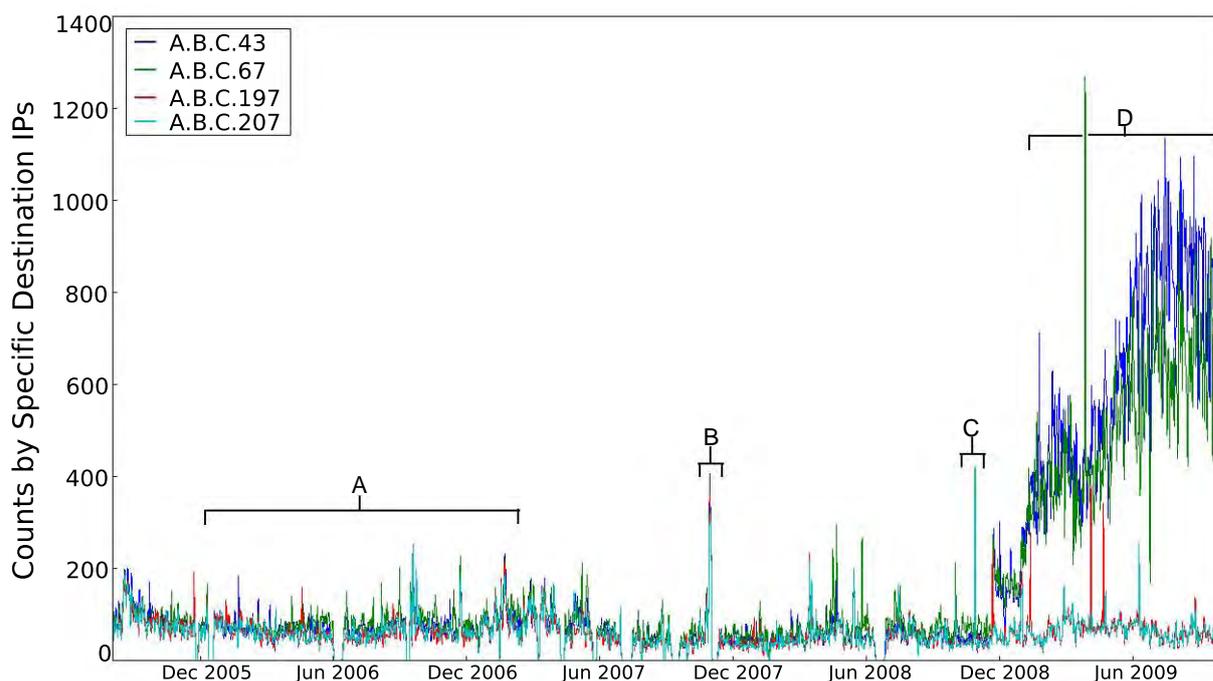


Figure 5.4: Packets Counts per Destination IP (/day).

from the four IPs tend to move together without any significant difference. However from point D onwards, the time span of Conficker, there appears to be a disconnection between IPs above and below a.b.c.128. This is clearly caused by the flaw in the Conficker host selection algorithm, as previously discussed in Section 3.4.1. While the case illustrated by point D is a clear indicator of malicious activity, it requires pre-knowledge that there was such a flaw in Conficker and thus picking which IPs of the network telescope IP space to monitor is challenging, especially for network telescopes which monitor large IP spaces.

#### 5.2.4 Port and Type Counts

Considering counts for specific ports and ICMP types has distinct value when identifying malicious incidents that may occur due to recent discovery of vulnerabilities in a specific protocol or application. Figure 5.5 illustrates the history of outbreaks associated with 445/tcp as seen by the network telescope. Points A-C illustrate large spikes that were not found in any of the surveyed literature or analysis of the Wireshark captures that could conclude any possible reasons for the increase in 445/tcp at these times. This highlights the both the weakness of passive telescopes and further that not every spike is necessarily identifiable as a type of specific behaviour or as a new trend. Point D is thought to be

related to Conficker pre-scanning that took place before the main outbreak of Conficker in order to identify vulnerable hosts for pre-seeding purposes. It is clear from point E that there is a clear and sharp increase in 445/tcp traffic which is of course consistent with the outbreak of Conficker. Again the issue of how to select which ports are worthy of analysis crops up, though this issue may be solved by considering which ports are currently topical according to recent trends.

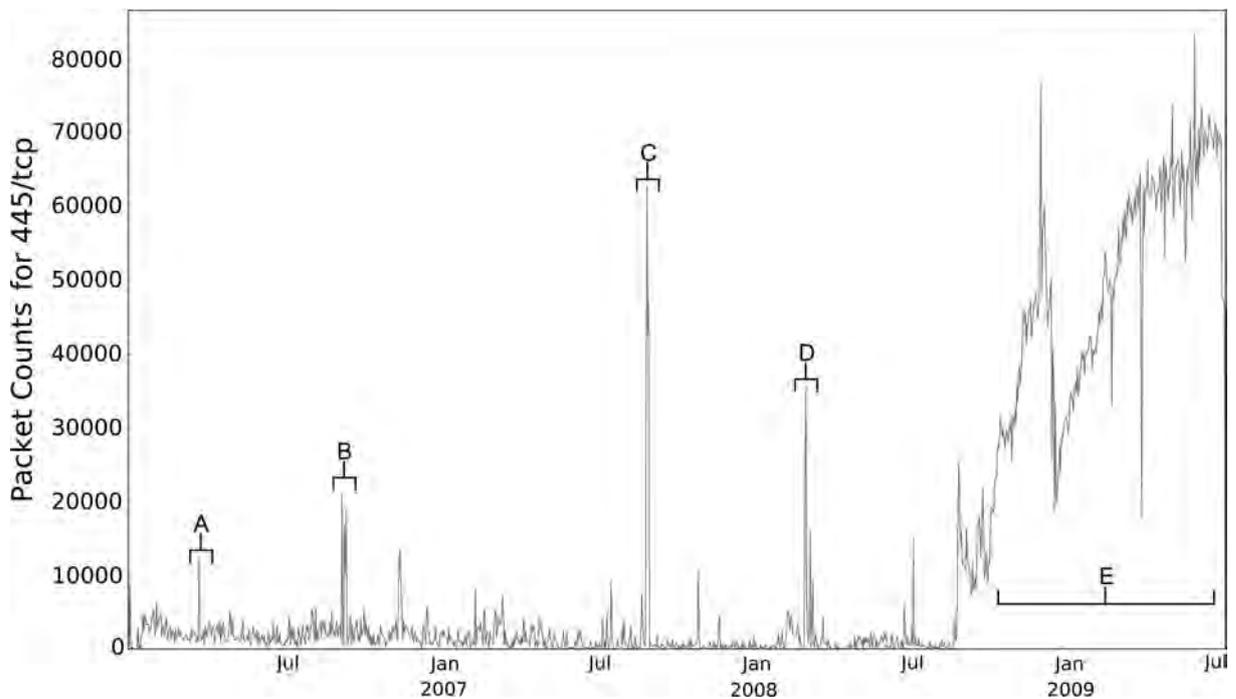


Figure 5.5: Packets Counts for Dest. Port 445/tcp (/day).

An alternative approach to studying specific ports is to instead monitor the top  $n$  ports or types for a particular time period. Figure 5.6 notes the changes in top five TCP ports per year. The yearly level was chosen as it would be difficult to visualize any such metric for shorter periods. It is clear from this figure that before Conficker the majority of traffic is made up by 'the other' ports. However there is a major shift between 2008 and 2009 with 445/tcp becoming the dominant destination port for traffic received by the network telescope. This suggests a shift in malicious activity that is worthy of investigation by an analyst.

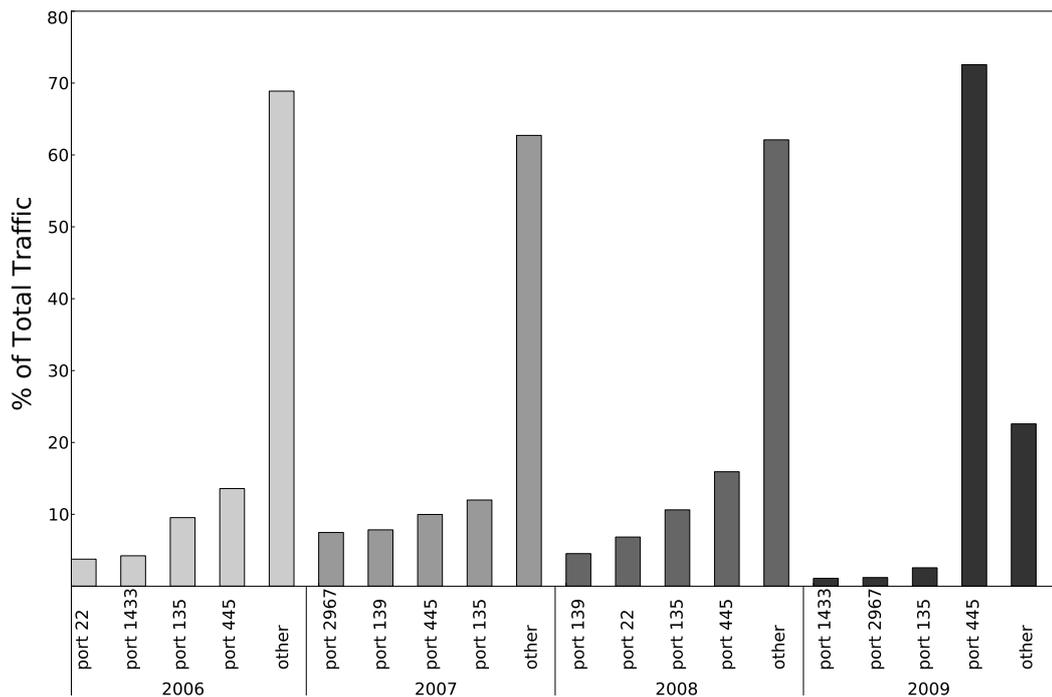


Figure 5.6: Top Five TCP Ports (/year).

## 5.3 Analysis by Aggregation

One of the major issue identified in the previous section was that a number of metrics are difficult to display due to large size of the input space. A solution to this issue is to aggregate these variables through functions such as average, max, minimum and variance. These techniques were chosen as they are well known and simple to implement though it is noted that it is possible that other aggregation techniques could be more effective. The following aggregated metrics are discussed in this chapter: average packet size, average TTL value, average counts per source IP, average counts per destination IP and maximum packets per destination IP.

### 5.3.1 Average Packet Size

With regards to changes in the average packet size metric it is important to consider that there is significance in both an increase or a decrease in packet size as either could suggest a major change in the composition of traffic. Although, having said that there are possible cases where the average may remain very similar or least not differ significantly from previous values as two incidents balance each other in deforming the average. An

important factor to remember with this indicator is the bounded nature of packet sizes, as discussed in Section 3.2, as this defines limits on the range of the average packets size. Figure 5.7 graphs the average packet size calculated per day.

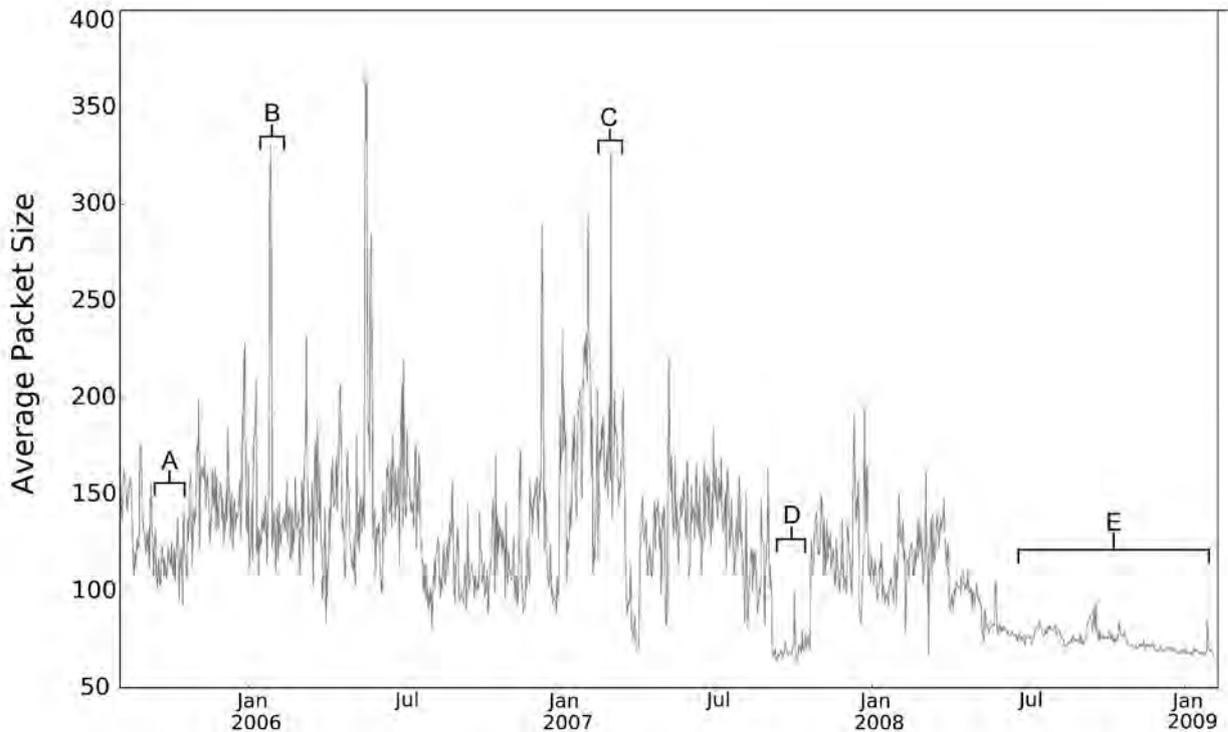


Figure 5.7: Average Packet Size (/day).

It is expected that individual cases of extreme behaviour will be reduced by increasing the size of the grouping from which to take the average. The average packet spikes upwards of 380 bytes and drops to lows of 80 bytes though for the majority of the observed period the average moves around 120 bytes. Point A shows an example of breadth scanning that targeted the network telescope for a number of days. Points B and C are both examples of net send malware which greatly inflate the average packet size as this attack tends to have unusually large packet sizes as was described in Section 4.3. Point D seems to be due to the interaction of scanning thought to be related to Conficker. Finally point E shows how a large scale event like Conficker affects the average packet size metric. Conficker pulls the average down greatly due to its high volume of small packets making it difficult to infer further activity from this metric.

### 5.3.2 Average TTL Values

The TTL values in the studied datasets, as was shown in Section 3.4.3, have three distinct modes which can be attributed to the characteristics of the major OS groups. Generally speaking when modes are identified within a dataset the average statistic becomes a less appropriate measure of central tendency. However, this thesis is interested in averages as a source of incident identification and not as an accurate statistical description of the data.

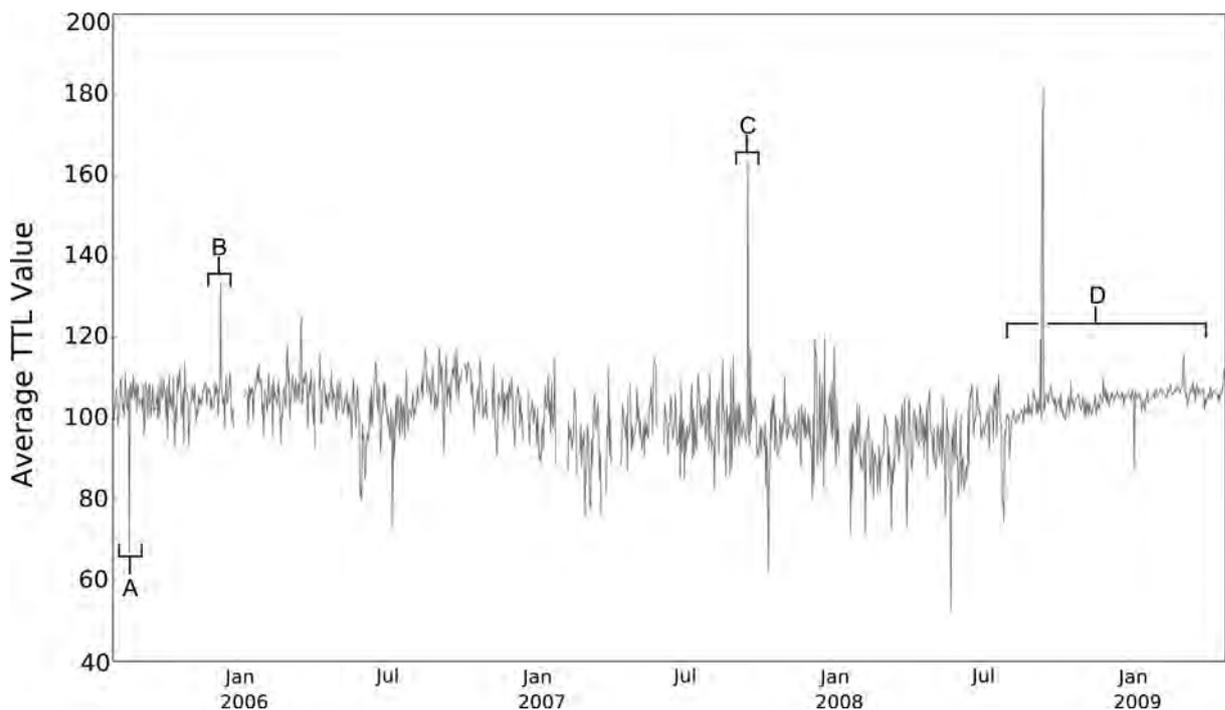


Figure 5.8: Average TTL Value (/day).

The average TTL metric, as displayed in Figure 5.8, keeps somewhere between a value of 80 and 120 though a number of large spikes and dips are noted. Points B and C are attributed to large DDoS attacks that took place, as DDoS attacks tend to be associated with hosts which fall in TTL Mode C, as was mentioned in Section 3.4.3. Compared to other metrics the average TTL is a stable due to the dominance of the Windows OSs as shown in Section 3.4.8 with the exception of a number of spikes and dips due to sudden floods. It noted that post Conficker, as shown by point E, the average TTL variability is narrower than previously observed which is attributed to the increased proportion of traffic that is received from Windows hosts.

### 5.3.3 Average Counts per Source IP

The average counts per source IP metric is simply the average number of packets received per source host in a given time window. It is expected that this metric will be of small magnitude as it is unlikely that hosts repeatedly communicate with IPs belonging to the network telescope as this is not in the best interest of worm authors or people scanning for vulnerabilities. This metric is expected to be able to identify cases where there is a dramatic change in the number of packets received by specific hosts or a general shift in the number of packets generally received per host.

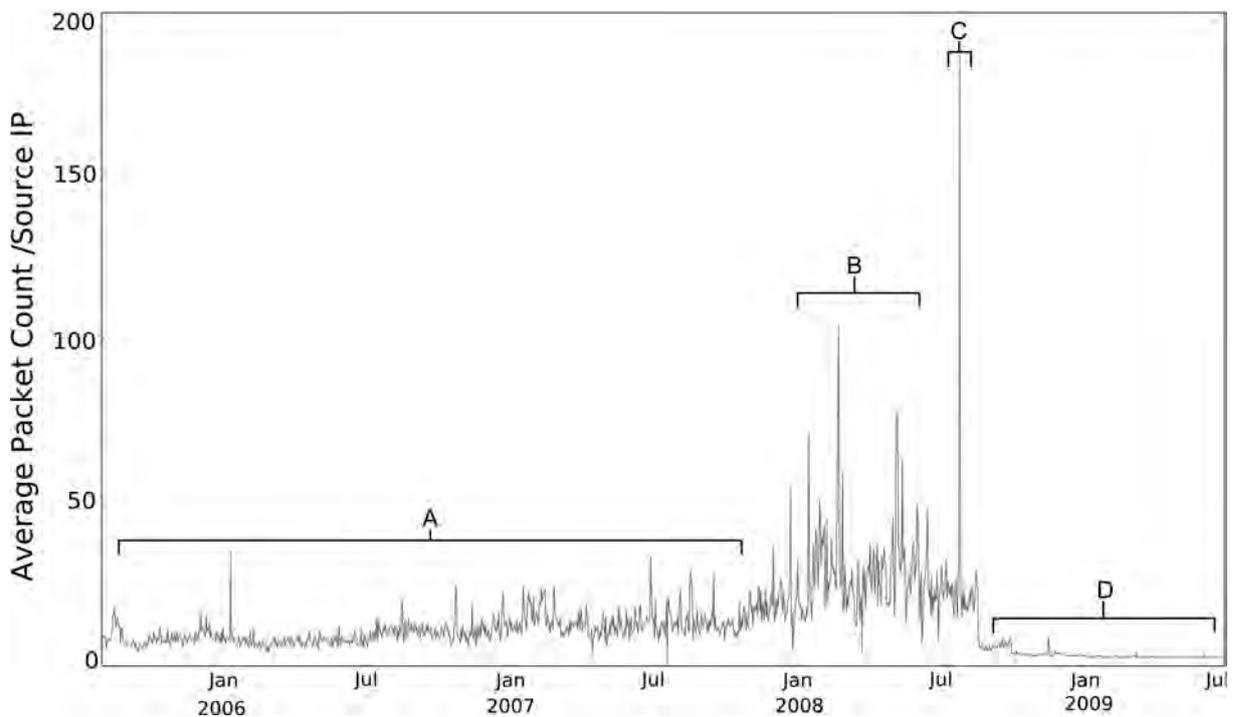


Figure 5.9: Average Packet Count per Source IP (/day).

Figure 5.9 plots the average counts per source IP per day as observed by the RU1 dataset. For period A the average packet count per Source IP remains relatively stable in a range of approximately [7-10] packets, which could be argued to be relatively low. A number of cases of high volume scanning and DDoS of a number of telescope IPs are noted in the spikes such as point C. During point B, Conficker pre-scanning, there is far higher spikes which is sensible as there are a number of source IPs. Point D illustrates how a change in the composition of traffic greatly affects many metrics, which is good initially for identifying unusual activity but becomes troublesome when trying to identify new malicious activity. Conficker reduces the average count per source IP to the range of [1-3]

packets per source IP.

### 5.3.4 Average Counts per Destination IP

The average number of packets received per destination IP bears a great amount of resemblance to the total packet counts metric which is somewhat expected although the magnitude of packets received is lower.

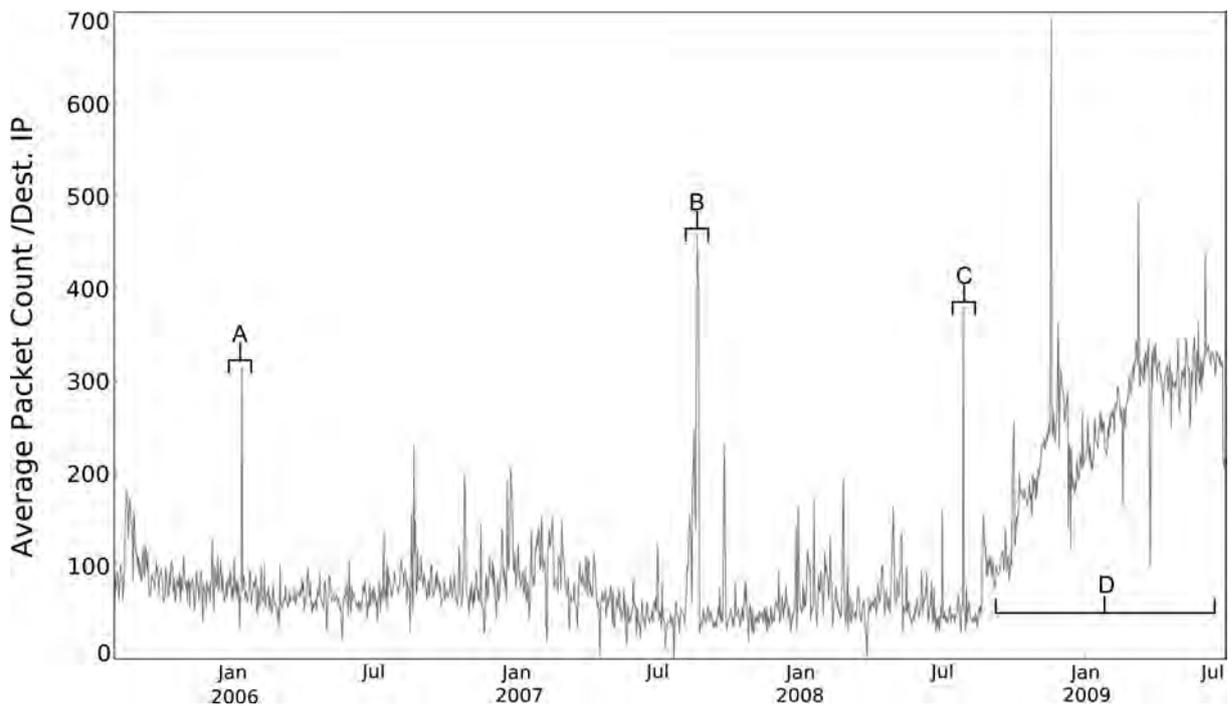


Figure 5.10: Average Packet Count per Destination IP (/day).

Figure 5.10 depicts the average counts per destination IP. For the most part, the average counts per destination IP metric hovers around 100 packet per day, spiking up towards 700 packets during cases of scanning and DDoS and then exhibits the characteristic logistic growth during the outbreak of Conficker, as shown by period D, as many other metrics have identified. Points A, B and C were investigated and found to be related to points of DDoS as these attacks greatly increasing the packet counts received for a few IPs causing a large shift in the average.

### 5.3.5 Maximum Packets per Destination IP

An issue with using averages is that it is possible to weigh down the extrema in the dataset which may be of interest to researchers. While certain incidents such as Conficker clearly have a significant effect on the nature of malicious activity and are thus clearly visible by many metrics, smaller incidents may still be identified by metrics which are better suited for this purpose such as the maximum packets per destination IP.

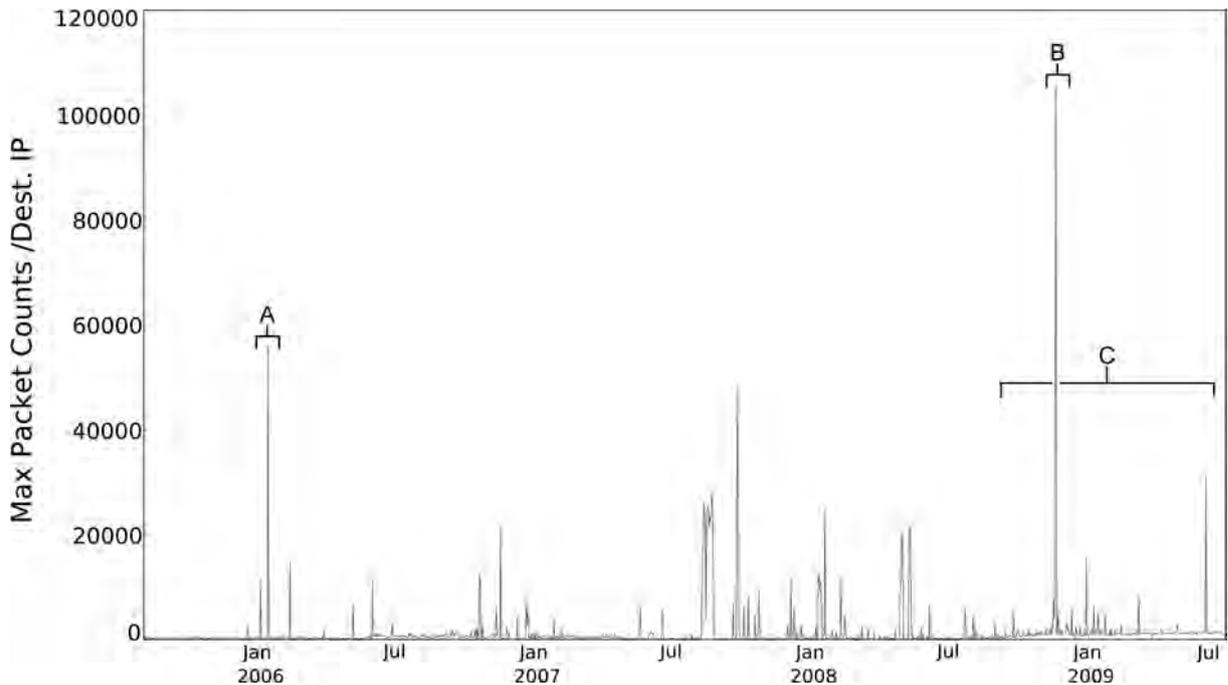


Figure 5.11: Maximum Packets per Destination IP (/day).

The maximum packets per destination IP calculated at a per day level is provided in Figure 5.11. This figure has drastically dissimilar behaviour to that found in many of the metrics examined so far, notably there is no significant trend during Conficker, as illustrated by point C, as observed in a number of other metrics. Points A and B identify cases of scanning and DDoS respectively as large spikes in the maximum number of packets observed per destination IP. It would thus seem there is some value for this metric if observing this behaviour easily was the intended goal.

## 5.4 Analysis by Groupings

An alternative approach to forming metrics that solves the issue of dealing with a large input size, as was the main issue identified in Section 5.2, is to group by a certain quantity such as the Operating Systems sub-group or by varying IP netblocks. In this section protocol type, OS groupings, TTL groupings, groupings by netblock and number of bogon IPs observed are studied.

### 5.4.1 Protocol Counts

Counting the number of packets received at a per transport protocol level allows for a similar holistic understanding of the data as is gleaned from the total packet counts but gives slightly deeper insight into what a particular incident may or may not be. For example, an increase in the amount of ICMP traffic is most likely some form of DDoS and is unlikely to some form of a worm. Figure 5.12 depicts the change in the counts per protocol at a daily level compared with the ratio of counts per protocol per day. The count per protocol mirrors the behaviour seen in the total packet count metric and this would be sensible as counts per protocol is merely a decomposition of the total packet count. It is clear that more activity is visible from this metric than the total packet count as there is less interaction amongst incidents as is illustrated by the fact that the DDoS that occurs at Point B is unobscured by the Conficker period shown during period C.

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. Through it needs to be considered whether an increase in the proportion of one quantity relative to another is actually an increase in the one, a decrease in the other quantity or a combination of both of these effects. As is expected TCP is the dominant proportion of total traffic followed by UDP and ICMP respectively, as shown by Figure 5.12. TCP makes upwards of 60% of the total traffic observed. UDP varies from approximately 0% to 40% of the total traffic and ICMP starts off at approximately 20% and varies around this value then spikes upwards of 30% at point B but as a general trend decreases with time. Point A illustrates the potential weakness when inferring the ratio as an identification mechanism as the ratio shows a decrease in TCP and an increase in UDP, while there was in fact an increase of UDP traffic at this time there was also an increase in TCP traffic, though this increase is not sufficient to balance the increase in UDP traffic to retain a

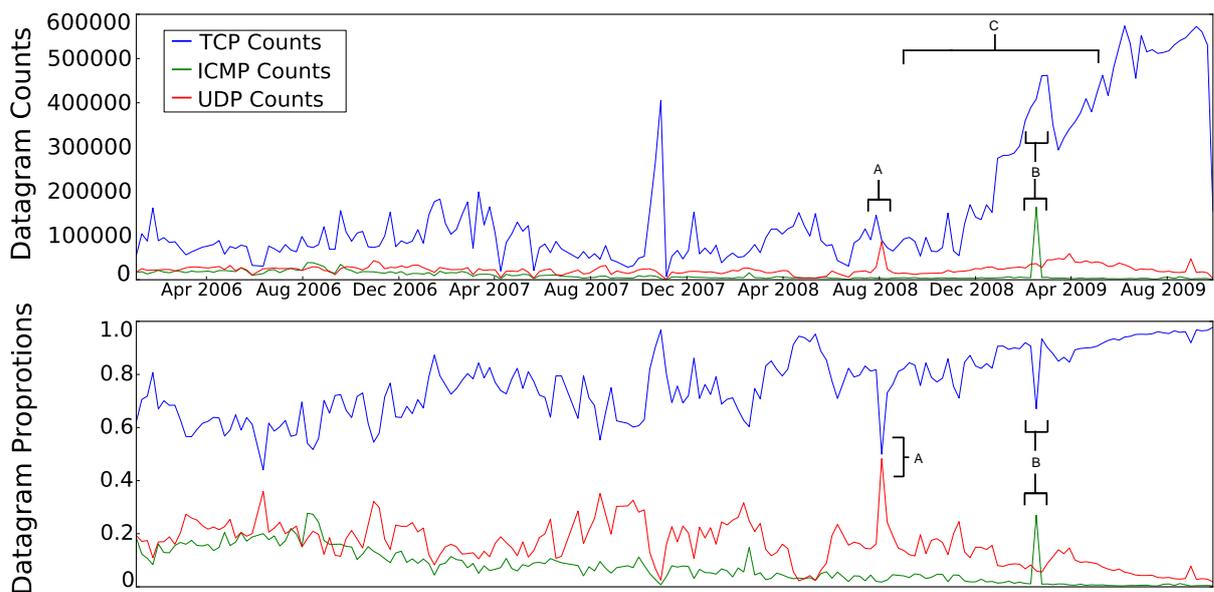


Figure 5.12: Packets Counts per Protocol (/day).

similar ratio.

## 5.4.2 OS Grouping

An approach may be to consider the packet counts per time period for groupings of the major OSs. Three simple groups could be devised: Windows and related platforms, Unix and BSD based platforms and a catch all of the remaining OS types. The OS groupings used are listed in Table 5.2, it is noted that the OS identification was achieved using p0f. Figure 5.13 is a line chart displaying the variation in these groupings over the time span of the RU1 dataset.

The Windows grouping shows behaviour similar to that of the total packet count and the TCP portion of the protocol count metrics. This is sensible as the dominant OS contributor will have a large impact on high level metrics as illustrated by period C, showing the characteristic logistic growth associated with Conficker. The Linux grouping has a much lower magnitude of traffic received but still highlights some points of interest via large spikes in the data as shown by points A and D. Point A is a case of DDoS undocumented in this thesis while point D is related to DDoS.D as discussed in Chapter 4. Finally, the ‘other’ grouping receives far less packets on a daily basis than Windows OS, but a number of minor spikes such as B are still noted. Point B was attributed to

net send malware.

| Windows Grouping           | Linux BSD Group   | Other Grouping       |
|----------------------------|-------------------|----------------------|
| Windows 2003 (AS)          | FreeBSD 2-6       | Cisco 12008          |
| Windows XP with SP2/SP3    | Linux 2.0.3 - 2.6 | Eagle Secure Gateway |
| Windows 2000 with SP1 - P4 | MacOS 9.0 - 9.2   | Redline T—X 2200     |
| Windows CE                 | OpenBSD 3.0 - 3.9 | HP-UX B.10.20        |
| Windows 98                 | Solaris 7 - 10    | PocketPC 2002        |
| Windows NT 4.0 SP6a        | NetBSD            | Proxyblocker         |
| Windows 95 (b)             | BSD 3.1 - 4.3     | NMAP scans           |

Table 5.2: Operating Systems Groupings.

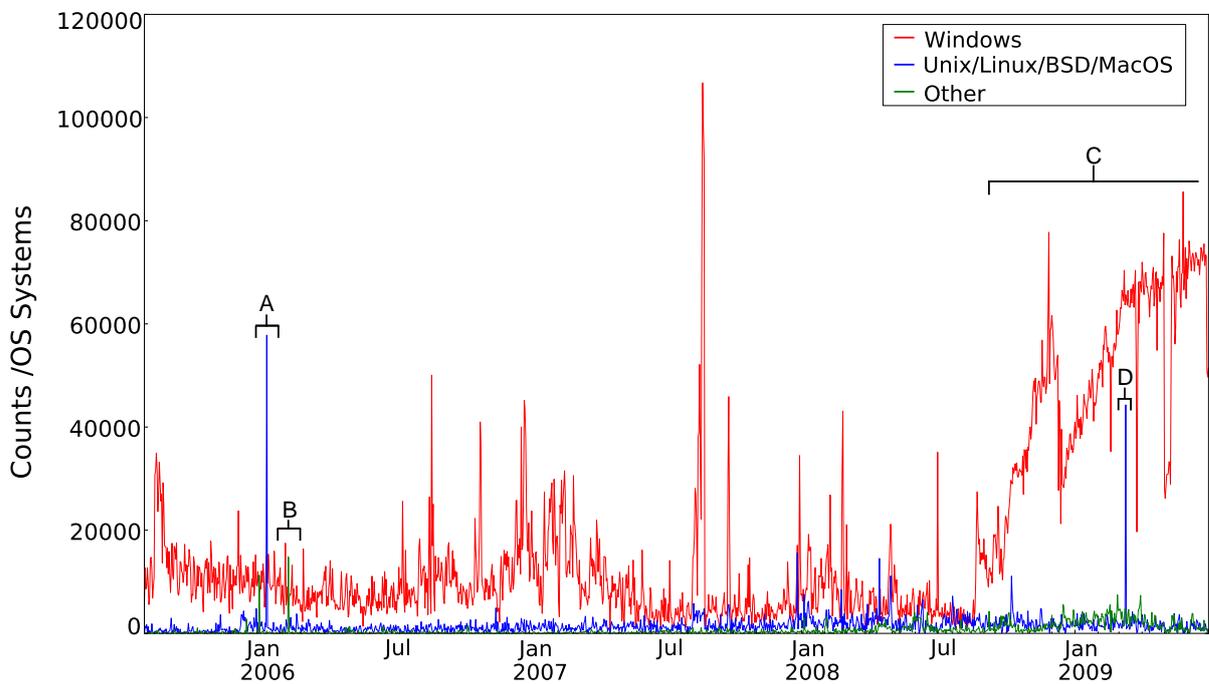


Figure 5.13: Packet Counts per OS Grouping (/day).

### 5.4.3 Bogon Counts

Counting packets received with a source IP in the bogon IP space is expected to illuminate cases of DDoS and spoofing as the individuals conducting these attacks tend to value their anonymity. Generally speaking, the bogon counts at a per day level is low relative to the total count as this metric hovers in the range [0-5] packets per day, as shown in Figure 5.14. However, there are a large number of significant spikes in the dataset, some reaching

packet counts in the region of 700/day. Point A was found to be related to a number of days of depth port scanning targeted towards the network telescope as is supported by the raising in the low bound of the number packets received /day shown in the figure. Points B and C are related to cases of hard depth scanning. During point D, the period associated with Conficker, the characteristic logistic shape of growth in infected Conficker hosts is again present through considerably down scaled.

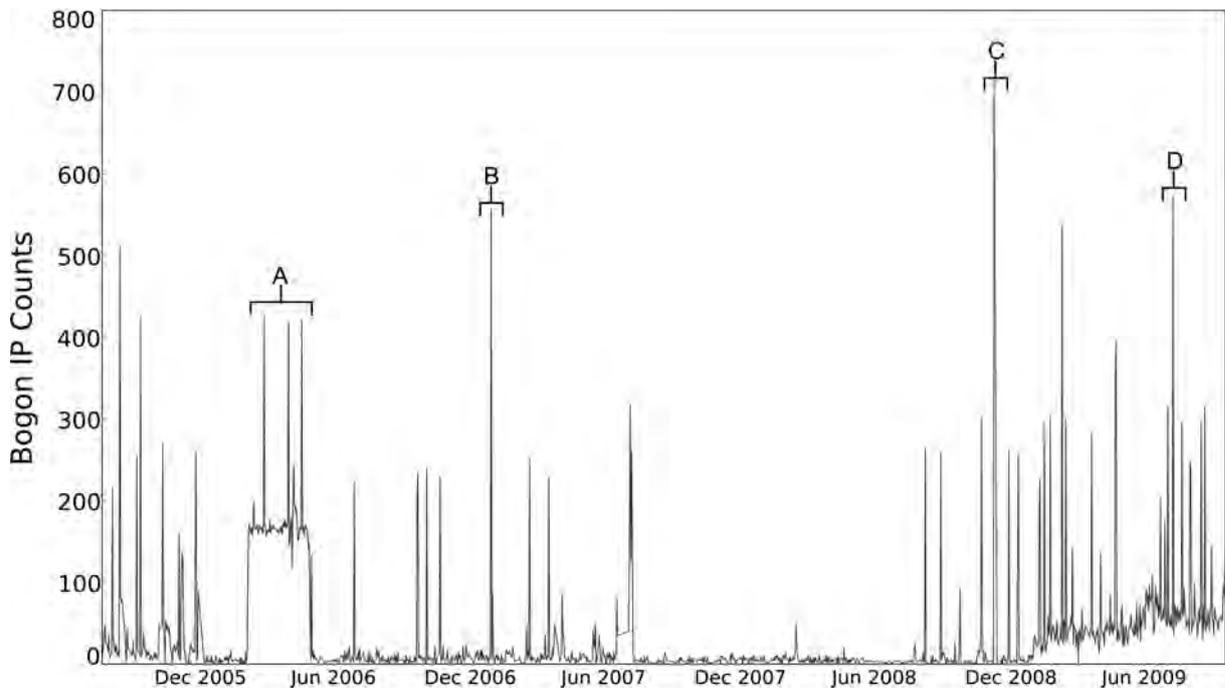


Figure 5.14: Bogon IP Counts (/day).

## 5.5 Summary

This section has taken an initial look at a number of measures that can be derived from network telescope data. Section 5.1 discusses the idea of normality with respect to network telescope metrics and provided an example by examining the change in total packet counts metric. In particular the affect that Conficker has on this metric was discussed as a violation of normality.

Section 5.2 considers a number of count-based metrics. These were defined as metrics that required no further calculation than simply enumerating a quantity by a given criteria. Counts calculated in this section included: total packet count, frequency of specific packet

sizes, frequency by destination IP and specific port or ICMP types. An important issue noted from the packet size, destination IP, ports and ICMP type counts is that the input space for each of these metrics is too large to consider the frequency of each particular value. For instance there is simply too many destination IPs for a typical telescope to be able to plot the number of packets each IP receives.

A sensible approach to reducing the amount of data and complexity in a dataset is to aggregate data in a specific and intelligent way. In this vein Section 5.3 lists and briefly evaluates aggregated metrics. Average packet size, average TTL, average counts per destination IP, average counts per source IP and maximum counts per destination IP were considered in this section. An issue identified in this section is how large incidents such as Conficker influence these aggregates making it difficult to infer future behaviour relative to historical values.

Section 5.4 looks at the formulation of metrics by grouping according to specific criteria. Metrics discussed in this section included: datagram counts, datagram ratio's, OS grouping and bogon counts. The datagrams and OS groupings exhibited behaviour similar to that observed by other metrics such as the total packet counts and counts by destination port, though in a more segmented and clear way. The number of packets that claimed to originate from bogon IP space was shown to be low in volume but identified cases of scanning and further showed a scaled down version of Conficker's growth which was consistent with metrics previously discussed

A number of potentially useful metrics have been identified and discussed in this chapter and have been shown to be able to identify particular malicious behaviour. While these metrics have shown to be useful it would be useful if this process of identification could be slightly more automated, it is not expected that a single system is capable of accurately identifying all incidents but rather tools should be investigated to help support analysts. In the next chapter baselines and technical analysis are considered as identification aids.

*History never repeats itself, but it often rhymes.*

Mark Twain.

# 6

## Incident Identification Techniques

**D**ETERMINING whether the change in a quantity is sufficiently significant to highlight a potential incident was one of the major concerns identified in Chapter 5. This is due to the fact that there is no way to be certain that a spike, dip or other unusual behaviour in a given quantity is truly anomalous without investigating the actual packet capture. As this process is lengthy and intensive for an analyst to perform it would be advantageous to reduce the number of cases that are to be investigated.

A way to alleviate this issue is to devise signalling based on metrics derived from network telescope datasets. This chapter considers techniques from baseline analysis and technical analysis. As it would not be feasible to discuss application of these techniques to every metric discussed in Chapter 5, a few examples are chosen to illustrate specific points. The remains of this section are ordered as follows:

- Section 6.1 applies three different types of baselines: specific-time baseline, recurring baselines and rolling baselines to a selection of metrics.
- Section 6.2 tests a number of the technical analysis techniques including Donchian price channels, Bollinger bands, Keltner channels and the moving average envelopes.

- Section 6.3 summarizes the topics discussed in this chapter.

The technical analysis section of this was formed from papers written for the SAICSIT (South African Institute for Computer Scientists and Information Technologists) 2010 Conference (Cowie, 2010a) and the SATNAC 2011 Conference (Cowie, 2010c).

## 6.1 Baselines Analysis

A baseline can be used in a network telescope analysis scenario to provide evidence to suggest that a major incident has occurred by citing that a particular measure was previously much higher or lower. Three types of baselines shall be considered within this thesis, namely: times-specific baselines, rolling baselines and recurring baselines. The baselines discussed will be formed as a percentage as this provides easier interpretation.

### 6.1.1 Specific-time Baselines

Specific-time baselines contrasts current observations with those made at a specific time period (Holman, Johnson and Bradley, 2009). An example of this is comparing the number of ICMP type 11 messages observed in the last 24 hours when compared with days where large DDoS attacks were observed.

$$B_{specific} = 100 \frac{C_{current}}{C_{specific}} \quad (6.1)$$

The specific baseline, as given by equation (6.1), as one may expect merely scales the original metric by an arbitrarily chosen constant. This may be useful if there are some properties of a particular metric that constraint it such that a particular value is significant. However there are no obvious properties for any of the metrics discussed in Chapter 5. Arguably the packet size property has significant values at the MTU and the minimum valid IP packet, creating baselines based on these properties does not provide any new information and may only be useful to identify cases where there is a strong dominance of poorly forged packets. The other major concern is that certain incidents may become difficult to observe, for example a small DDoS attack may be completely unobservable from this baseline when compared to a very large DDoS. Due to these concerns the

specific-time baseline was not investigated further in this thesis, though an example may be found on the accompanying CD.

### 6.1.2 Recurring Baselines

Recurring baselines, as shown in equation (6.2) compare current readings against those made in the same period (Holman, Johnson and Bradley, 2009). Analysing the differences between metrics taken between Q1 2009, Q1 2010 and Q1 2011 provides an example of a recurring baseline.

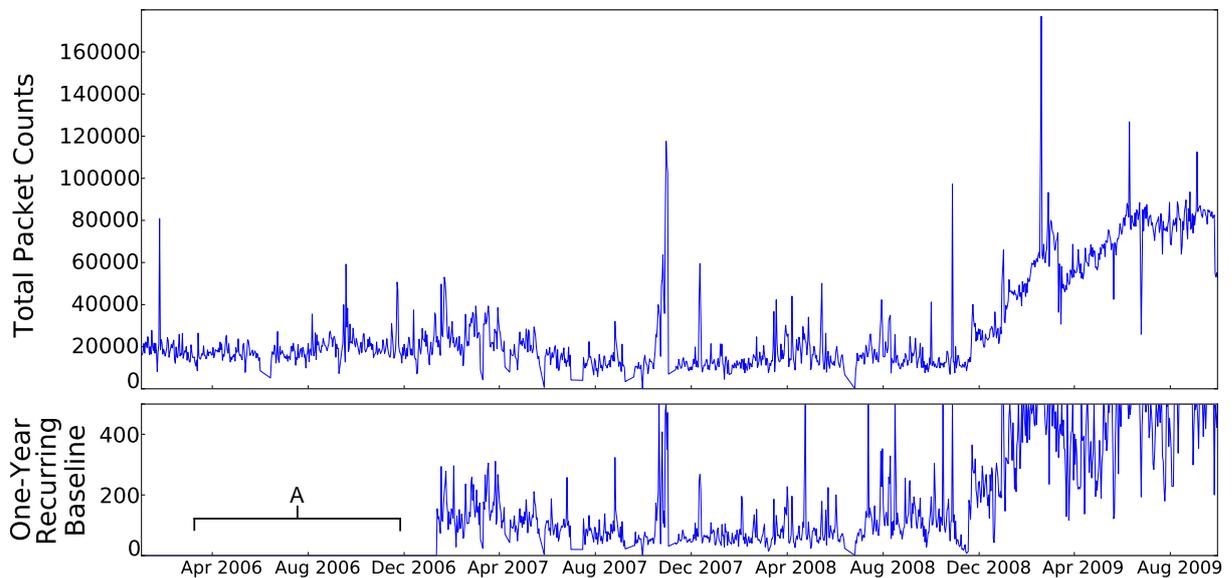


Figure 6.1: One-Year Recurring Baseline of Total Traffic (/day).

$$B_{recurring} = 100 \frac{C_{current}}{C_{current-n}} \quad (6.2)$$

The major concern with this type of baseline is what is the significant relation between a value measured today and  $n$  periods ago. As a result there may be a significant difference between two points. Further, it is noted that for periods before  $n$ , this baseline is undefined. It is possible that there may have been a large spike in traffic at the previous point while the current point is relatively normal resulting in a significant dip in the baseline value. This is illustrated by period A, as the difference between the points at A and one-year previous is so great, it overshadows the other values in the dataset. As such this

approach needs to be applied to metrics that remain relatively stable or that have bounds such as average packet size as soon in Point B. One solution may be to limit the maximum percentage difference as shown in row two of figure or make use of a logarithmic scale.

### 6.1.3 Rolling Baselines

Rolling baselines make a comparison between current measurements and a statistical calculation made over the last  $n$  periods (Holman, Johnson and Bradley, 2009). For example, the number of TCP packets counted on a specific day could be compared to the average packet count recorded for the last six months. Equation (6.3) gives a mathematical formulation of a rolling baseline.

$$B_{rolling} = 100 \frac{C_{current}}{Mean(C_{previous})} \quad (6.3)$$

A number of possible period could be chosen to form a rolling baseline; section investigates the behaviour of a six-month baseline formed from the total packet count metric as is displayed in Figure 6.2

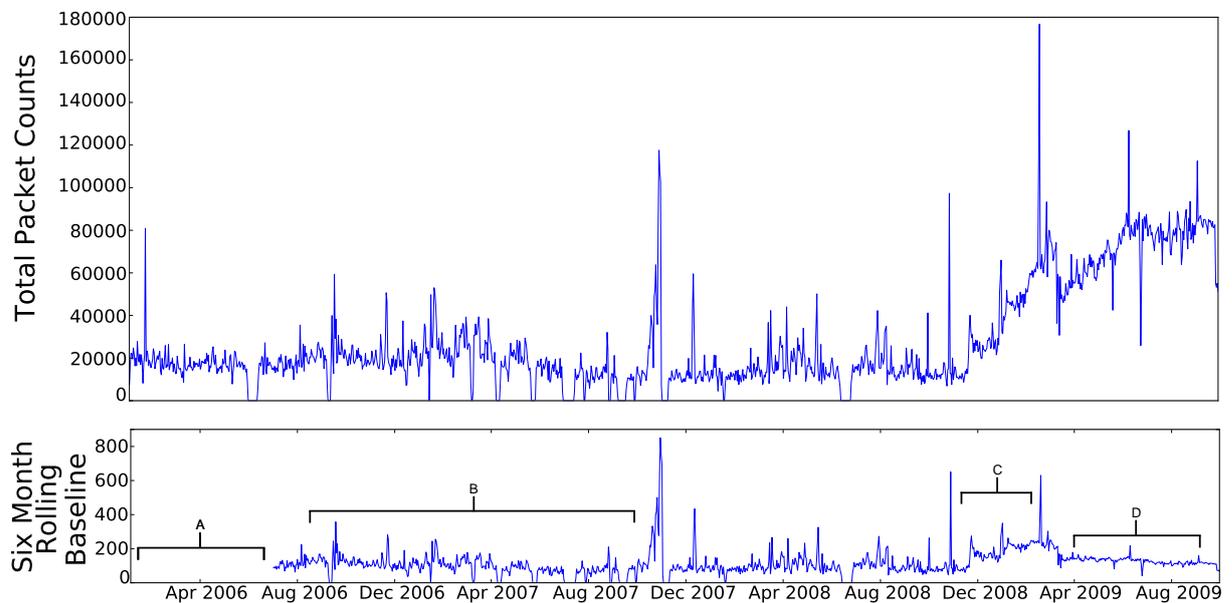


Figure 6.2: Six-Month Rolling Average Baseline for Total Traffic. (/day)

It is worth noting that a rolling baseline will lag by the period of the recurring baseline

as shown by period A. In this case a rather large period was chosen and as such a large section of the timeline will have no accompanying signal. The baseline shares a similar shape to the original metric with similar spikes and dips, as shown by periods B and D. Large extrema as is exhibited by point C do seem to maintained as long as this behaviour does not become the new norm. As for interpreting the baseline as a signal it appears that anything above a 200% appears to relate to major incidents and identifies a number of incidents discussed in Chapter 4, though it is a somewhat arbitrary rule. Point C, the logistic growth of Conficker, does show some growth in the rolling baseline but it not as distinct as is found for the original metric, the previously lower counts act as a counter balance to lower the signal observed. Finally, Point E shows the point from which the growth in Conficker is less rampant, as a result the baseline does not seem to indicate any behaviour that would suggest investigation solely from this baseline.

The maximum statistic is an alternative aggregation technique that may form a useful for baseline for metrics that have tend to made up of large spikes instead of a general such as ICMP type 11. Figure 6.3 shows the one-month maximum rolling baseline. In comparison to the specific-time baseline for the same metric, period A shows that there is more value to the rolling maximum baseline as there is movement in the baseline during this period from which activity could be inferred.

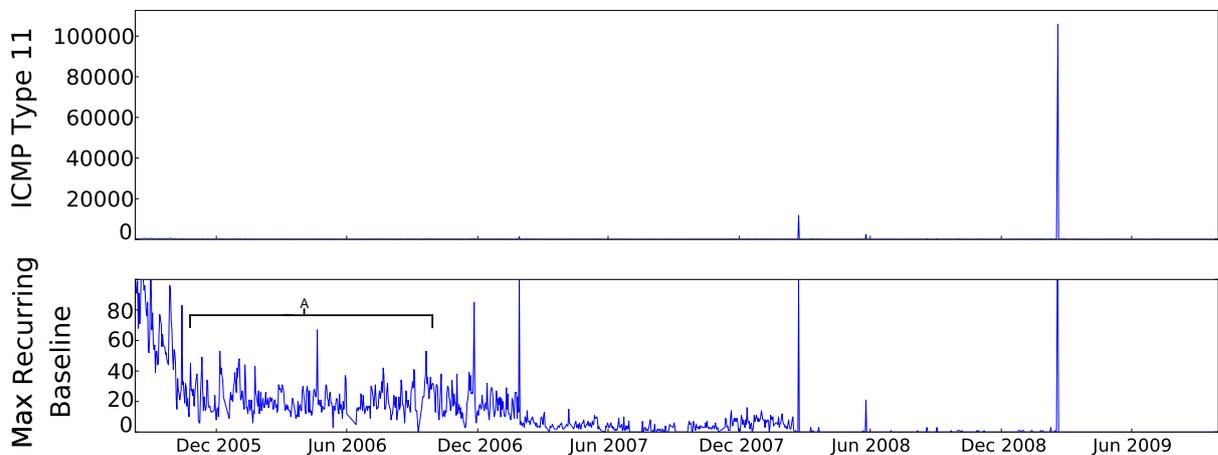


Figure 6.3: One-Month Rolling Maximum Baseline for ICMP type 11 .

## 6.2 Technical Analysis

An interesting approach is to consider network telescope data as if malicious traffic is a market with trends and quantities and then to apply techniques from technical analysis to find changes in the ‘market’. A selection of technical analysis indicators and trading bands are considered in this section, including: Donchian price channels, Bollinger bands, Keltner channels, moving average convergence divergence and moving average envelopes.

### 6.2.1 Donchian Price Channels

Price channels define a simple band through a set of three lines, one above the quantity, one below the quantity and the final being the midpoint between the other two lines which are formed from maximum and minimum values for a given period. Let  $D_n$  denote a list of the last  $n$  values for the quantity then equations (6.4), (6.5) and (6.6) define the bands of a standard Price Channel

$$PC_{middle} = \frac{\min(D_n) + \max(D_n)}{2} \quad (6.4)$$

$$PC_{lower} = \min(D_n) \quad (6.5)$$

$$PC_{upper} = \max(D_n) \quad (6.6)$$

The typical approach to trading when using price channels is to observe when the commodities price rises above the previous high and then buy. When the prices drops below the previous low the trader should sell (Chart School, 2010). This is commonly referred to as the  $n$ -day rule (technical market indicator, pg 421). A more concrete version of this rule is Donchian’s 4-Week Rule, this effectively means a period of 20 trading days. For a first attempt at applying price channels to the data set, the 4-week rule shall be applied. A strategy that could be derived is to say if a cross over occurs wait until the opposite cross over occurs and define that as a point of interest worth investigating. The rationale behind this is for cases where an incident occurs for a few time periods, brute force scanning for example may only last a few hours or in extreme cases a few days causing a drastic reduction in the average packet size. An alternative may be to consider that a large increase in a short space of time may be indicative of erratic behaviour which may be caused by increasingly larger bursts of traffic, that could be from worm outbreaks.

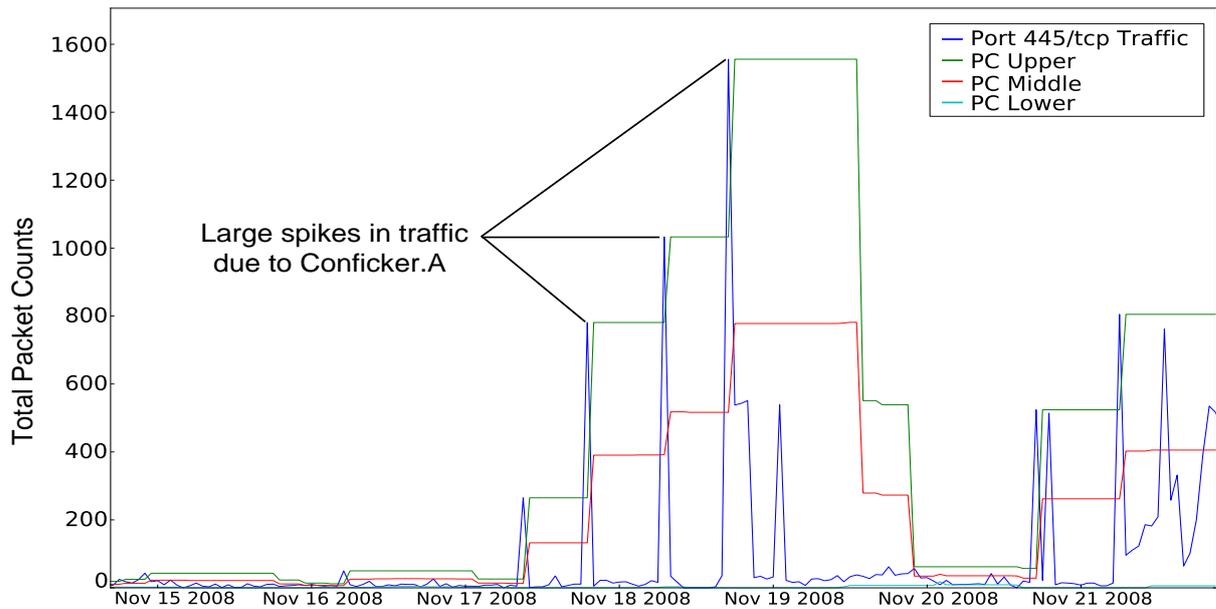


Figure 6.4: Price Channels from 445/tcp Dest. Port Packet Counts (Conficker.A) (/day).

Figure 6.4 considers price channels at the point in time of the outbreak of Conficker A. The three significant spikes thought to be a surge in 445/tcp scanning are identified by a crossing over of the quantity by the upper band. The sharp nature of the bands changing is clear from this image. This property is somewhat desirable in that it works well to generate a signal if a large spike is followed by smaller spikes as only one signal is generated. However, in cases where a particular metric grows uniformly with time a large amount of signalling will occur.

Figure 6.5 applies price channels to a case of TCP port scanning, as noted by period A. The average packet count barely dips below the bottom band of the price channel three times during period A. As the incident period is sufficiently long the upper band falls down as there is no longer a large value in the last 20 entries. The upper band is then crossed over at two points by the average packet size during the incident period. This behaviour is somewhat undesirable, it would be hoped that a crossing of the upper band would signal the end of such a period however this example shows that this is not true for a given period. Activity before point A is also of interest. At an hourly level the average packet size tends to vary greatly due to the interactions of the various types of malicious activity pulling the average packet size up and down. As a result a number of cross-overs occur before the period (11 times crossing above, 11 drops below). This makes it difficult to able to differentiate points of interest as a number of signals would have been registered in a relatively short period.

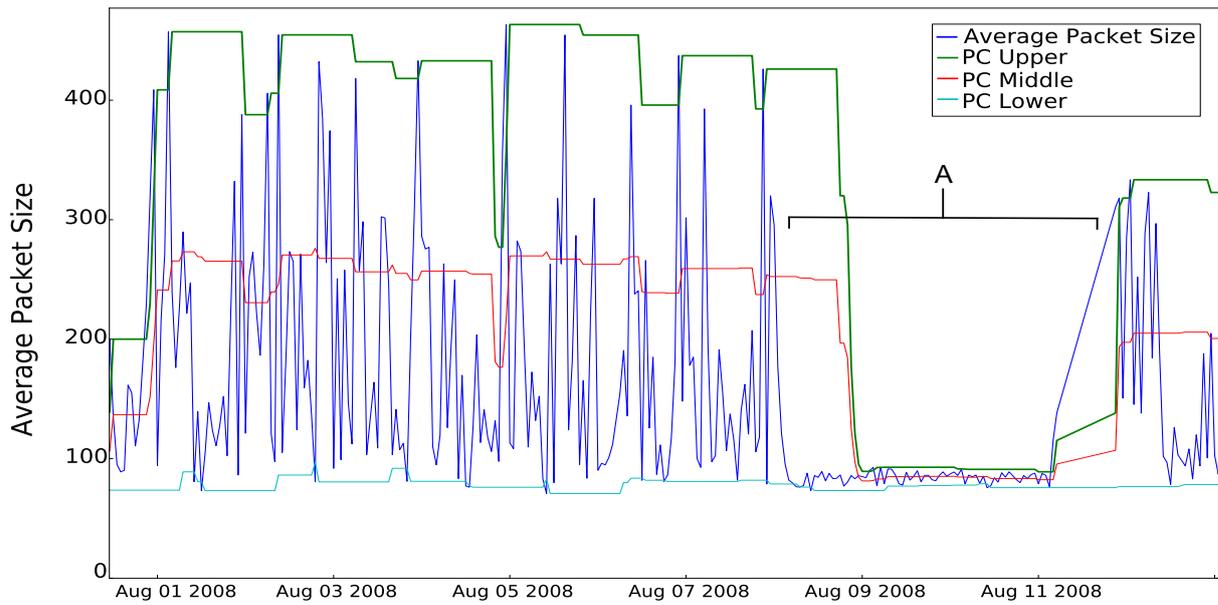


Figure 6.5: Price Channels from Average Packet Size (Scan.A) (/day).

### 6.2.2 Bollinger Bands

Bollinger bands are a popular technical analysis tool that allow for a relative definition of the highness or lowness of a quantity as compared to previous values. Bollinger bands consist of two main components, a measure of central tendency together with a measure of volatility (Bollinger, 2002a). The commonly accepted measure of central tendency is a moving average of a quantity while the volatility is usually expressed by the standard deviation of said quantity. These two components yield three bands commonly known as the upper, middle and lower Bollinger bands. These bands are defined by equations (6.7), (6.8) and (6.9).

$$B_{lower} = MA(n, d) - 2\sigma \quad (6.7)$$

$$B_{middle} = MA(n, d) \quad (6.8)$$

$$B_{upper} = MA(n, d) + 2\sigma \quad (6.9)$$

One of the indicators that can be derived from Bollinger bands is Bollinger percent band (%b) as defined by equation (6.10). The %b indicator describes the current observation in relation to the upper and lower bands. When %b exceeds one the observation is above

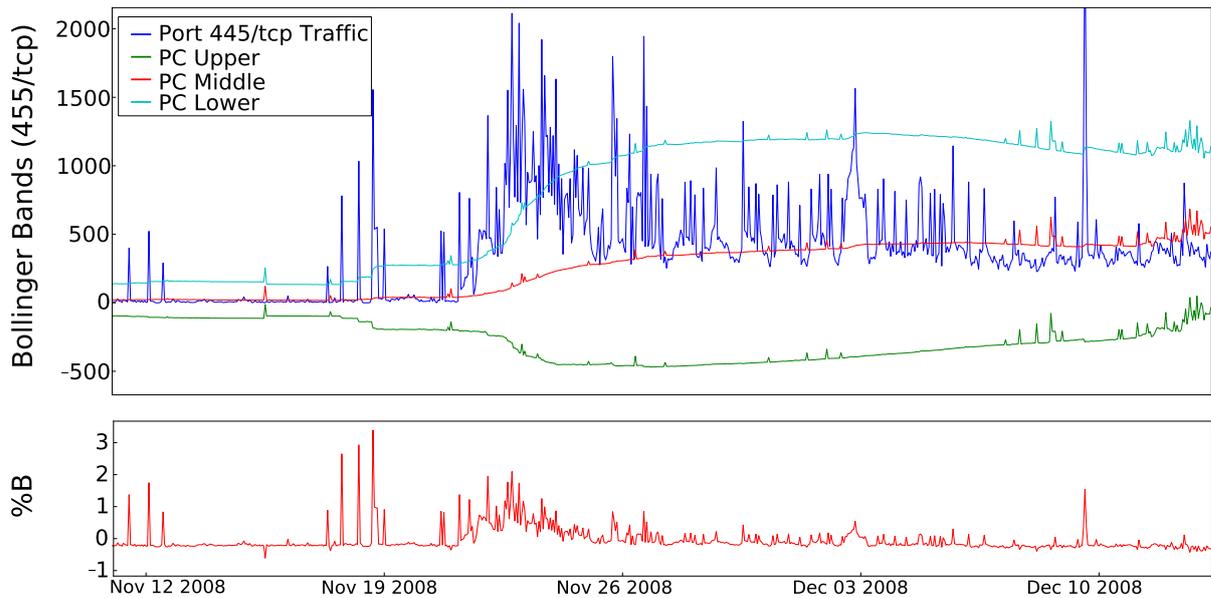


Figure 6.6: Bollinger Bands (%b) from 445/tcp Dest. Port Packet Counts (/day).

the upper band and when it is less than minus one it is below the lower band (Bollinger, 2002c).

$$\%b = \frac{\text{last} - \text{lower}BB}{\text{upper}BB - \text{lower}BB} \quad (6.10)$$

### %b as a Signal

When the value of %b is greater than one, the quantity has crossed over the upper Bollinger band. Many traders use this as a signal to indicate that the quantity is oversold. While there is no directly analogous concept of oversold within the field of network telescopes the concept of crossing the band is still useful as illustrated in Figure 6.6. The %b indicator becomes greater than one a total of three times during the breakout time of Conficker.A and thus identifies it strongly.

### 6.2.3 Keltner Channels

Keltner channels make use of a moving average for the measure of the central tendency and the Average True Range (ATR) provides a measure of the volatility. ATR is based upon the True Range value. True Range for a given day is defined as being the largest of the current high for the day less the current low for the day, the absolute value of the high of the day less the previous close or the absolute value of the most recent periods low less the previous close (Wilder, 1978). The Average True Range is then calculated as given by equation (6.11).

$$ATR_n = \frac{ATR_{n-1} + TR_n}{14} \quad (6.11)$$

The Keltner Channel may then be calculated using equations (6.12) and (6.13).

$$KC_{lower} = MA(n, d) - k * ATR \quad (6.12)$$

$$KC_{upper} = MA(n, d) + k * ATR \quad (6.13)$$

The moving average to be considered in this paper will be a 20 day SMA and  $k = 2$ . The simplistic rule of trading when the quantity crosses the channel lines is still applicable and will be the strategy of choice (Colby, 2002). Figure 6.7 shows sample application of Keltner Channels to network telescope data. It is clear that the packet counts for 445/tcp clearly cross over the upper Keltner Channels at about the time Conficker.A emerged, though the distance between the bands remains the same.

### 6.2.4 Moving Average Envelopes

Moving average envelopes are constructed from a moving average such as an EMA or SMA from which the upper and lower bounds are  $\alpha\%$  percent below or above. Due to the percentage based nature of the upper and lower bands it follows that bands will most likely be rather rigid and slow to change (Murphy, 1999a). Equation (6.14) gives the mathematical form of these lines.

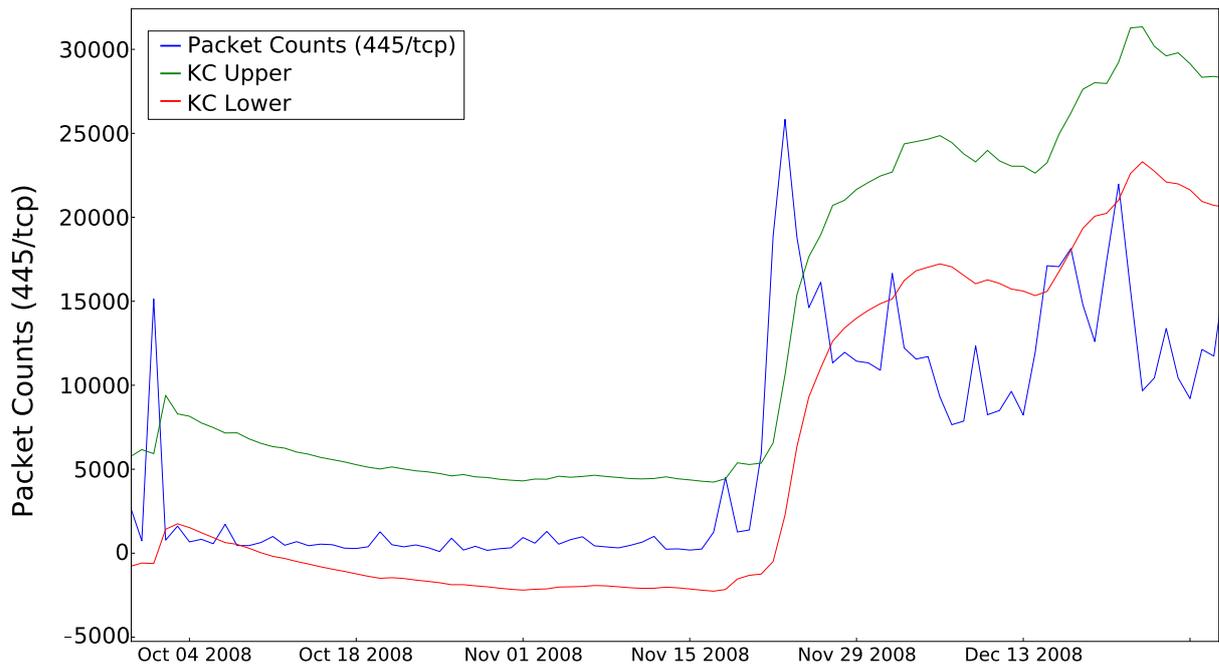


Figure 6.7: Keltner Channels from 445/tcp Dest. Port Packet Counts (/day).

$$\begin{aligned}
 E_{upper} &= (1 + \alpha)SMA(x, n) \\
 E_{lower} &= (1 - \alpha)SMA(x, n)
 \end{aligned}
 \tag{6.14}$$

A sample application of moving average envelopes to port 445/tcp is provided in Figure 6.8. Pre-Conficker, the bands are very close together as shown by period A but the bands adjust slowly due to the sudden increase caused by point B results in an expansion of the. Though this expansion is slow and another crossover occurs again due to the large spike shown at point C. The %b indicator borrowed from Bollinger bands identifies these points of cross-over.

## 6.3 Summary

This chapter has detailed a number of approaches that can be applied to metrics to help identify points of interest within a network telescope dataset. The chapter began by considering types of baselines and how they may be applied to network telescope datasets. Rolling baselines, recurring-time baselines and specific date baselines were discussed and

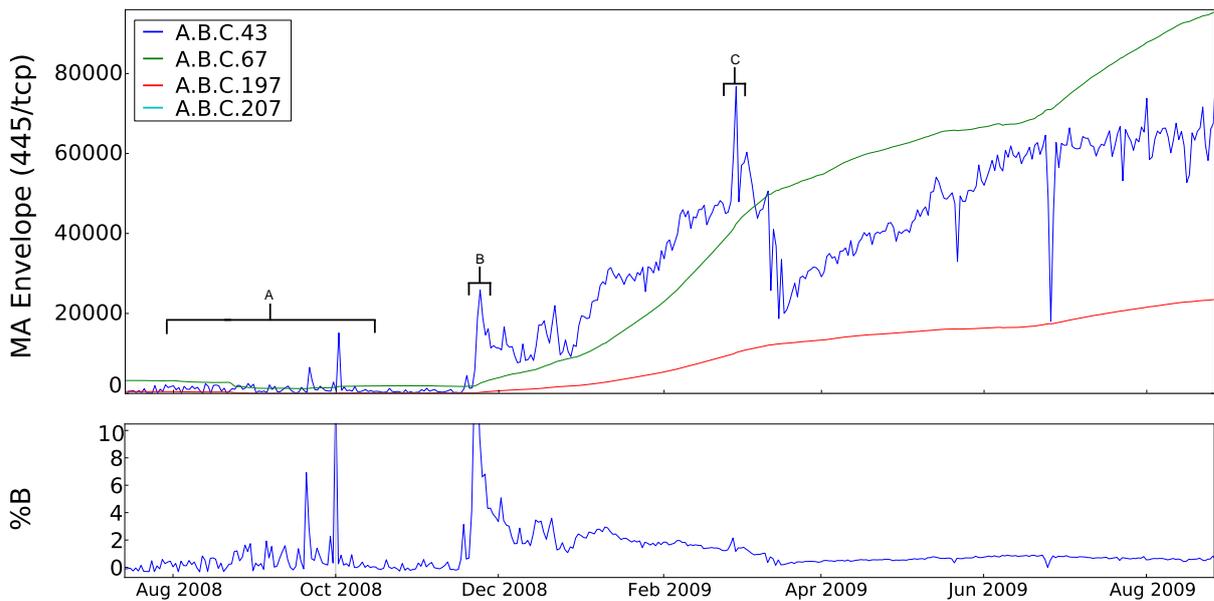


Figure 6.8: Moving Average Envelopes from 445/tcp Dest. Port Packet Counts (/day).

applied to metrics in Section 6.1. It was shown that specific-time and recurring baselines are not particularly effective though in some specialist case they could be of use. Rolling baselines

Technical analysis indicators and bands are considered in Section 6.2 as a way to highlight new trends in malicious traffic. The TA indicators considered in this section included price channels, Bollinger bands, Keltner channels and exponential moving envelopes. Price channels were shown to effectively identify both the logistic growth of Conficker and the drop in average packet size due to scanning. Bollinger bands were investigated by application to Confickers growth which showed that the bands can identify this incident and readjust appropriately so that excessive amounts of over signalling did not occur. As a result the Keltner channels suffered from issues relating to a poor mapping from original scenario that they were used in when placed in a network telescope context as they function on the values high, low and close from the financial context. As close has no real significance in the network telescope scenario and the low may often be 0 for a given quantity it follows that these indicators are often badly affected and often nonsensical. Moving average envelopes were shown as an alternative formation of a band which while being successful at identifying Conficker seems to suffer from bands that adjust too slowly.

Having concluded the metrics and analysis section of this thesis the next chapter considers application of the approaches and metrics to incidents identified in Chapter 4.

*However beautiful the strategy, you should occasionally look at the results.*

Winston Churchill

# 7

## Results

**T**HIS chapter assesses the metrics and techniques discussed in Chapters 5 and 6 by evaluating how they react to incidents described in Chapter 4. The metrics evaluated were categorized as providing either strong identification, weak identification or failure to identify by considering whether they produced a visual signal that could be used to identify the incident. Strong identification implies that the signal generated is clear and unambiguous. Weak identification implies that the signal hints at the possibility of a potential incident but the signal generated is not clear. Failed was defined as showing no change that would alert an observer to a potential incident.

Metrics which successfully identified specific behaviour were then used as input data to the identification techniques. The results from the identification techniques were categorized as having identified the incident or having failed to identify the incident. It is noted that it is insufficient to suggest that a metric or identification technique that has more success at identifying incidents is superior to other approaches as this may exclude the issues associated with said metric due to the limited nature of this study. As such the particular strengths and weakness of these approaches are included where possible.

This chapter is divided into two main sections and a summary section as listed below:

- Section 7.1 considers the effectiveness of metrics formed by counts, aggregation and grouping as outlined in Chapter 5.
- Section 7.2 evaluates baseline analysis and technical analysis as discussed in Chapter 6.
- Section 7.3 notes the recommendations of the researchers with respect to which metrics and identification approaches show promise.
- Section 7.4 concludes this chapter by summarizing the major points discussed in this chapter.

## 7.1 Metrics

A number of metrics were discussed in Chapter 5 by considering three approaches to metric formation, namely: count-based metrics, metrics formed through aggregation and grouping based metrics. These approaches are now evaluated by considering the major incidents observed in Chapter 4.

### 7.1.1 Count-Based Metrics

The major advantage of count-based metrics is that they are fast to calculate as they only require simple mathematical and data organization operations and as a result tend to be simple to interpret. In this section the total packet count, counts per packet size, counts per specific destination IP, counts per specific destination port or ICMP type are considered in terms of their ability to identify changes in malicious traffic. Table 7.1 lists the ability of count-based metrics to observe the behaviour detailed in Chapter 4.

The total packets received metric as shown in Table 7.1 often fails to identify incidents of a smaller scale as shown by the lack of significant identification provided for the minor DDoS cases, net send malware and scanning cases. However, for the significantly larger incidents, such as Conficker, the total packet count metric does show sufficient evidence of change to suggest the need for investigation.

| Incident           | Metrics      |             |                  |                 |                |        |        |
|--------------------|--------------|-------------|------------------|-----------------|----------------|--------|--------|
|                    | Total Counts | Size Counts | Port/Type Counts | Top Ports/Types | Dest IP counts |        |        |
| DDoS A             | Failed       | Failed      | Strong           | Strong          | Failed         | Strong | Failed |
| DDoS B             | Failed       | Failed      | Strong           | Strong          | Failed         | Failed | Failed |
| DDoS C             | Failed       | Failed      | Weak             | Weak            | Failed         | Failed | Failed |
| DDoS D             | Failed       | Failed      | Strong           | Strong          | Failed         | Failed | Failed |
| DDoS E             | Failed       | Failed      | Failed           | Failed          | Failed         | Failed | Failed |
| Conficker A        | Strong       | Strong      | Strong           | Strong          | Strong         | Strong | Failed |
| Conficker B        | Weak         | Strong      | Weak             | Strong          | Failed         | Strong | Failed |
| Conficker C        | Strong       | Strong      | Weak             | Weak            | Strong         | Strong | Failed |
| Conficker D        | Failed       | Strong      | Weak             | Weak            | Strong         | Strong | Failed |
| Conficker E        | Failed       | Strong      | Weak             | Weak            | Strong         | Strong | Failed |
| Net Send Malware A | Weak         | Strong      | Strong           | Strong          | Strong         | Weak   | Failed |
| Net Send Malware B | Weak         | Strong      | Strong           | Strong          | Failed         | Failed | Failed |
| Net Send Malware C | Failed       | Strong      | Strong           | Strong          | Failed         | Failed | Failed |
| Net Send Malware D | Failed       | Strong      | Strong           | Strong          | Failed         | Failed | Failed |
| Scanning A         | Failed       | Strong      | Strong           | Strong          | Strong         | Failed | Strong |
| Scanning B         | Failed       | Strong      | Strong           | Strong          | Failed         | Failed | Weak   |
| Scanning C         | Failed       | Strong      | Weak             | Weak            | Failed         | Failed | Failed |
| Scanning D         | Failed       | Strong      | Weak             | Weak            | Failed         | Failed | Weak   |

Table 7.1: Identification Ability of Count-Based Metrics.

The core weakness of the total packet count metric is that it only implies that there may be some significant change in the composition of malicious traffic but it does not suggest the finer details required for incident identification, though arguably if logistic growth is observed it is likely that some form of worm is responsible. The total packet counts metric seems appropriate as a metric that should be a component of a network telescope dashboards or analysis frameworks as it provides a quick overall glance at the state of the malicious traffic on the Internet.

Counting by specific packet sizes can be of some use when identifying unusual behaviour. Counting specific packet sizes identifies Conficker strongly when counting 62 byte packets specifically and cases of net send malware, when the correct size is chosen, but both of these cases require pre-knowledge. Having specific graphs for packets of exactly 564, 168 and 455 bytes, as shown in Chapter 5, is too specific to be of use for general day-to-day analysis of incoming telescope data. It may be of use to analysts to pick a number of the more popular packet sizes, such as 62 bytes, and monitor the proportion that these sizes make of the total traffic.

Counting by specific source and destination IPs is challenging without any pre-knowledge of what IPs are likely to generate malicious traffic due to the large input space. Further, picking a few choice IPs appears to be an ineffective way to identify unusual activity, as shown in Table 7.1. As was mentioned in Section 3.4.1 the disparity in packets received above and below a.b.c.127 was noted and acts as a good indicator of Conficker. For cases of scanning, if one observes the individual hosts that are scanned it is clear that some activity has occurred but this would involve monitoring each host which is infeasible for large network telescopes.

Specific port and ICMP type counts show success at identifying all forms of activity, if one knows what activity is to be identified. For example, observing DDoS by considering ICMP type 11 reveals evidence of all of the considered DDoS attacks, monitoring 445/tcp clearly shows evidence of Conficker, monitoring 1026-1027/udp shows evidence of net send malware. As was the case with packet sizes, there are simply too many ports to keep track of each port and thus a top ports metric does seem to be the best way to approach this type of analysis

Monitoring the change in top  $n$  ports or types is effective at identifying the macro changes in the composition of malicious activity. This is shown by how the metric noted the strong increase in 445/tcp traffic accompanied by the outbreak of Conficker. However, many of the more minor incidents such as scanning, net send malware and DDoS simply do not

have a sufficiently large affect on the composition of malicious activity. The problem with such a metric is deciding how to visualize it on a periodic basis and further that this sort of measure needs to be taken for a fairly large period of measurement. That is to say, measuring changes in top  $n$  ports on a per hour basis is most likely too volatile a time period as sudden floods of traffic would greatly skew this metric.

This study of a selection of count-based metrics has shown that ‘simple’ measures can be effective for identifying a selection of incidents. Only DDoS.C was unidentified through count-based metrics and this could be attributed to the low packet counts associated with this incident. The major issues identified by the researchers from this section is the difficulty in selecting which counts to select for analysis particularly for packet size, TTL, specific ports and ICMP types without some knowledge that there is an incident related to a particular value of one of these properties. This study has only selected six possible simple counts, many were not explored due to space constraints but can be found on the accompanying C. A plethora of other count-based metrics could be devised from simple counts, with a few listed below: One weakness that may be expected with counts is, as the number of susceptible nodes increase on the Internet, comparison to historical values becomes less informative.

- Counts per specific source IP counts.
- Counts per specific OS types.
- Counts per geographical location.
- Counts per specific netblocks.
- Counts per specific Internet Registry

### 7.1.2 Aggregated Metrics

The input space for some of the variables studied were found to be too large to analyse each state individually, such as the packet size variable which had 809 different sizes observed in the RU1 dataset. These spaces can be reduced through the use of averages and extrema which were chosen as they have a clear interpretation and are relatively quick to calculate.

| Metrics            |              |               |           |          |                       |        |
|--------------------|--------------|---------------|-----------|----------|-----------------------|--------|
| Incident           | Ave. Src. IP | Ave. Dest. IP | Ave. Size | Ave. TTL | Max Packets / Dest IP |        |
| DDoS A             | Weak         | Weak          | Failed    | None     |                       | Strong |
| DDoS B             | Failed       | Failed        | Strong    | Minor    |                       | Strong |
| DDoS C             | Failed       | Failed        | Failed    | Strong   |                       | Failed |
| DDoS D             | Weak         | Strong        | Failed    | None     |                       | Weak   |
| DDoS E             | Failed       | Failed        | Failed    | Strong   |                       | Strong |
| Conficker A        | Strong       | Weak          | Failed    | Minor    |                       | Strong |
| Conficker B        | Weak         | Strong        | Failed    | None     |                       | Strong |
| Conficker C        | Major        | Strong        | Failed    | Strong   |                       | Failed |
| Conficker D        | Failed       | Strong        | Weak      | None     |                       | Failed |
| Conficker E        | Failed       | Strong        | Weak      | None     |                       | Weak   |
| Net Send Malware A | Weak         | Failed        | Strong    | None     |                       | Failed |
| Net Send Malware B | Failed       | Failed        | Strong    | None     |                       | Weak   |
| Net Send Malware C | Failed       | Failed        | Strong    | Minor    |                       | Failed |
| Net Send Malware D | Failed       | Failed        | Weak      | None     |                       | Failed |
| Scanning A         | Failed       | Failed        | Strong    | Strong   |                       | Failed |
| Scanning B         | Failed       | Failed        | Failed    | None     |                       | Strong |
| Scanning C         | Failed       | Strong        | Weak      | Minor    |                       | Strong |
| Scanning D         | Weak         | Weak          | Failed    | Minor    |                       | Weak   |

Table 7.2: Identification Ability of Aggregated Metrics.

This section discusses results for the following aggregated metrics: average counts for specific source IP, average counts for specific destination IP, average counts for TTL, average packet size and the maximum packet counts per destination IP. Table 7.2 summarizes the identification ability of these metrics.

The average source IP metric fails to identify incidents which do not define a major change in the composition of traffic traversing the Internet. As indicated by its failure to identify cases of DDoS, net send malware and scanning. This is due to the fact that individual attacks from a small number of hosts is unlikely to have enough of a measurable effect to cause a change in the average source IP. However, it is very capable of identifying the initial change from pre-Conficker outbreak and the Conficker outbreak.

The average destination IP appears to effective at identifying behaviour where there is a general rise in the packets per destination IP as shown by the identification of Conficker and cases of breadth scanning. This is further supported by the lack of evidence to identify some cases of DDoS and net send malware. It is worth noting again that the space studied by these telescopes is incredibly small so it is quite possible for a single depth scan to be able to affect this metric though this is unlikely for a large telescope such as those used by CAIDA.

The average packet size metric is interesting in that it can be influenced very heavily by the emergence of new behaviour as illustrated by its reaction to the outbreak of Conficker, as discussed in Section 4.2. As the packet size is bounded, the degree to which outliers can influence the metric is somewhat reduced. Large changes in the composition of malicious traffic which are identifiable by a change in packet size are generally observable from the average packet size metric with the exception of those incidents which have too little flow such as some of the cases of DDoS and scanning. There are a number of other aggregating functions which were not explored in this thesis that may be of interest as listed below:

Averaging the TTL value received would appear to have moderate success with identifying incidents. As TTL can be linked to OS type and the logical distance covered by the packet, attacks change either of these are detectable by the average TTL as illustrated by the identification of Conficker. The results for TTL do not appear to define a particular strength for the metric as it does observe DDoS.C, DDoS.E, Conficker.C and all the variants of net send malware.

The maximum packet count per destination IP correctly identifies the outbreak of Conficker.A, Conficker.B and identifies Scanning.C with minor identification of a case of DDoS

and Scanning.D. A number of cases of DDoS and scanning appear as spikes in specific metrics such as the total packet counts metric. It seems sensible that measuring the maximum would identify these cases, though the issue comes at choosing the length of the time period. By observation, taking the maximum packet count at a per day level is fairly easy to read and provides good signalling for DDoS and scanning.

It appears that aggregated metrics have some merit as a way to reduce the complexity of identifying malicious activity in a number of variables, all with large input spaces. Averaged quantities tend provide a good overall sense of the observed trend which is effective when dealing with incidents that completely alter the general composition of malicious activity. As evident by identification of Conficker by the average TTL, average packet size, average source and destination IP. A list of other aggregation techniques that could be employed is provided below.

- Alternative means such as windorized, trimmed and geometric means.
- Percentiles of the variables
- Medians of the variables
- Deviation of the other variables
- Maximum of the other variables

### 7.1.3 Grouped Metrics

Metrics that group variables by some criteria appear to be an effective mechanism for reducing the complexity of the network telescope problem space. This section discusses the results for: OS groupings, TTL Groupings, protocol groupings, protocol ratio's and bogon counts. Table 7.3 summarizes the ability of grouping metrics to identify incidents discussed in this thesis.

The OS grouping metric is useful in identifying cases where there is a shift in the composition of malicious activity due to the platform associated with said traffic. Though the core issue to consider that the Windows OS is an extremely dominant grouping to be begin with and as such the other grouping are going to hard to observe due to scaling. The Windows grouping does show a significant change during Conficker which makes sense due to the nature of the susceptible population

| Metrics.           |             |                 |        |
|--------------------|-------------|-----------------|--------|
| Incident           | OS Grouping | Protocol Counts | Bogons |
| DDoS A             | Failed      | Failed          | Failed |
| DDoS B             | Weak        | Failed          | Failed |
| DDoS C             | Failed      | Failed          | Failed |
| DDoS D             | Weak        | Failed          | Failed |
| DDoS E             | Failed      | Strong          | Failed |
| Conficker A        | Weak        | Strong          | Strong |
| Conficker B        | Weak        | Strong          | Failed |
| Conficker C        | Strong      | Strong          | Strong |
| Conficker D        | Failed      | Strong          | Failed |
| Conficker E        | Failed      | Strong          | Failed |
| Net Send Malware A | Failed      | Strong          | Weak   |
| Net Send Malware B | Failed      | Strong          | Failed |
| Net Send Malware C | Failed      | Strong          | Failed |
| Net Send Malware D | Failed      | Strong          | Failed |
| Scanning A         | Weak        | Strong          | Weak   |
| Scanning B         | Failed      | Strong          | Failed |
| Scanning C         | Strong      | Failed          | Failed |
| Scanning D         | Weak        | Weak            | Failed |

Table 7.3: Identification Ability of Grouping Metrics.

The ‘Other’ grouping was found to associate well with cases of scanning as expected while the Linux based grouping identified a number of cases of DDoS and net send malware.

Grouping by TCP, ICMP and UDP provides a high level analysis of the malicious traffic in transit of the Internet. This grouping was capable of identifying one of the large cases of DDoS, the Conficker variants, net send malware and a case of UDP scanning. This break down is also rational to analysts as a sudden increase in ICMP is most likely to be related to DDoS or ping flooding while UDP and TCP increase are most likely related to some form of scanning or worm activity. Through an analyst cannot be sure from this high level analysis what the particular type of activity is exactly without further inquiry.

Measuring the amount of traffic that ‘originates’ from bogon IP space does not appear to be an effective identifier of malicious activity, at least for the parameters of this study. The bogon counts metric failed to identify the majority malicious activity outlined in Chapter 4 and in general very little traffic from bogon IPs was received from both network telescopes. Though, as was mentioned in Chapter 5, the characteristic logistic growth of Conficker is observed by this metric. It is conceded that this may be specific to the network telescopes used and as such does not preclude bogon counts from being a useful metric, just one that is unfounded by this study.

Creating metrics from specific groupings allows for a more targeted approach to identifying incidents by highlighting a specific characteristic that is an associated characteristic of a potential incident type. A short list of other groupings that could be considered is presented below:

- Unique source IPs grouped at the /24,/16 and /8 subnets.
- Counts per continent.
- Counts per Internet registry.
- Counts per AS.

## 7.2 Identification Aids

This section considers the effectiveness of identification techniques to identify malicious activity in network telescope data. It would make little sense to apply analysis techniques

to metrics which failed to identify the malicious activity in the first place. As such, Table 7.4 states which metrics were applied to specific incidents. These metrics were chosen as they showed the best evidence of an incident having taken place, in the opinion of the researchers.

| <b>Incident</b>    | <b>Metric</b>                 |
|--------------------|-------------------------------|
| DDoS A             | ICMP type 11                  |
| DDoS B             | ICMP type 11                  |
| DDoS C             | Average TTL                   |
| DDoS D             | ICMP type 11                  |
| DDoS E             | Average TTL                   |
| Conficker A        | Total Counts                  |
| Conficker B        | Total Counts                  |
| Conficker C        | Total Counts                  |
| Conficker D        | Average Destination IP Counts |
| Conficker E        | Average Destination IP Counts |
| Net Send Malware A | Average Packet Size           |
| Net Send Malware B | Average Packet Size           |
| Net Send Malware C | Average Packet Size           |
| Net Send Malware D | 1026-1027/UDP Counts          |
| Scanning A         | Ave. TTL                      |
| Scanning B         | OS Grouping                   |
| Scanning C         | Max Packets /Dest IP          |
| Scanning D         | Size Counts                   |

Table 7.4: Metrics Chosen for Identification Aids.

### 7.2.1 Baseline Analysis

This section evaluates the baselines defined in Section 6.1, namely: recurring, six-month average and one-month maximum baselines. The specific-time baseline is not considered in this section as the researcher could not find significant dates to attributed to each metric. The recurring baseline did not produce any signal that was significantly different from the original metric. This adds no value to identification.

Taking a six-month rolling average baseline of the total packet count metric yielded some interesting and useful behaviour in that the percentage baseline does not radically shift due a major shift in a quantity. This is true as found for the large spikes in ICMP type 11 and counts for 445/tcp.

| Incident           | Metrics            |                           |                       |
|--------------------|--------------------|---------------------------|-----------------------|
|                    | One-Year Recurring | Six-Month Rolling Average | One Month Rolling Max |
| DDoS A             | Yes                | Yes                       | Yes                   |
| DDoS B             | Yes                | Yes                       | Yes                   |
| DDoS C             | Yes                | Yes                       | Yes                   |
| DDoS D             | Yes                | Yes                       | Yes                   |
| DDoS E             | Yes                | Yes                       | Yes                   |
| Conficker A        | Yes                | Yes                       | Yes                   |
| Conficker B        | Yes                | No                        | Yes                   |
| Conficker C        | Yes                | No                        | Yes                   |
| Conficker D        | Yes                | No                        | Yes                   |
| Conficker E        | Yes                | No                        | Yes                   |
| Net Send Malware A | Yes                | Yes                       | Yes                   |
| Net Send Malware B | Yes                | Yes                       | Yes                   |
| Net Send Malware C | Yes                | Yes                       | Yes                   |
| Net Send Malware D | Yes                | No                        | Yes                   |
| Scanning A         | Yes                | Yes                       | Yes                   |
| Scanning B         | Yes                | Yes                       | Yes                   |
| Scanning C         | Yes                | Yes                       | Yes                   |
| Scanning D         | Yes                | Yes                       | Yes                   |

Table 7.5: Identification Ability of Baseline Analysis.

However, for metrics where more stable behaviour is present, such as average TTL and average packet size, the six-month baseline contributes little to the original metric. It appears that rolling percentage baselines have merit as a means of considering current data without having to make a human based judgement on the relative size of extrema in the dataset. Further, the behaviour shown during Conficker encourages the notion that the six-month rolling baseline is still readable during a major incident. The one-year recurring baseline produces an image that is for many metrics such as total packet count or the ICMP type 11.

The one-month rolling maximum baseline does produce a useful signal for quantities that are dominated by generally low counts but have cases of large spikes such as the ICMP type 11 data as illustrated by the detection of DDoS.A,B and C. However, for metrics where the overall metric is not heavily dominated by extreme values, the rolling maximum mirrors the actual metric closely and thus exhibits the same detection ability as the actual metric. This is the case for net send malware and scanning.

In conclusion the six-month rolling average and one-month rolling maximum baselines do appear to be useful to analysis while recurring and specific-time baselines are more difficult to justify as useful to an analysis. Though the rolling maximum baseline does suggest that care needs to be taken when selecting an aggregating function for such a baseline.

## 7.2.2 Technical Analysis

This section evaluates the effectiveness of the following technical analysis indicators: price channels, Bollinger bands, Keltner channels and moving average envelopes. The success or failure of these techniques is listed in Table 7.6.

Price channels are a very simple calculation to both perform and to analyse and it would appear from Table 7.6 that they are successful in identifying unusual activity. Price channels, like all of the trading bands suffer, greatly from over signalling due to the very spiky nature of network telescope as opposed to the smoother nature of economic data. For price channels this over signalling appears to be especially strong due to the quick rate at which the bands contract or expand. If price channels are combined together with the simple cross over rule, all five cases of Conficker and DDoS are successfully identified. The bands react immediately whenever the rule comes into place but do not re-adjust until either the loopback period no longer contains the outlier or a new extrema is found

and thus bands tend to be rigid. This is a useful property for identifying incidents that occur on fairly smooth data. Price channels have the most success in identifying incidents however they also suffered from the most over signalling. Price Channels could be seen as a measure of support and resistance within the field of network telescopes.

Bollinger bands are far more volatile than price channels as the volatility is directly related to the deviation in the quantity. This property is useful when examining metrics that are extremely spiky such as counts for 445/tcp and 1433/tcp. The %b indicator provides a simple and easy to automate strategy to identify potential incidents. As the bands are more reactive, the amount of over signalling occurs is far less when compared to Price Channels. It is noted that %b failed to identify some of the Conficker variants. This is due to the fact that %b is extremely effective at identifying large spikes in the dataset but smaller spikes such as Conficker.C do not represent a significant enough change in previous readings to be registered as an incident.

The bands of the Keltner channel appear to be far less volatile than Bollinger Bands. As a visual analysis tool Keltner Channels tend to take a while to re-adjust after a major spike in traffic. The bands seem to remain the same distance apart until large spikes are present in the dataset. Keltner channels identified the majority of the incidents successfully and suffered minimal over signalling for the Conficker based incidents. However, for metrics where the 'low' of a particular period could be zero such as for ICMP type 11 counts the Keltner channels show unusual behaviour and provide no useful reading.

As moving average envelopes are percentage based they suffer greatly from over signalling due to the reaction of the bands. Moving average envelopes were capable of identifying a number of the cases of DDoS and net send malware occurrences but were unable to identify some of the later instances of Conficker. The issue here is the fact that the  $\alpha$  value used for moving envelopes is fixed so there is no context awareness, though it is noted that a SMA was used in this thesis and more complex moving averages may remediate this issue.

In conclusion, considering technical analysis as a means of identifying unusual activity has some merit but is fundamentally flawed. The assumptions of technical analysis, as stated in Section 2.4.4, do not map particularly well into the context of network telescope data. While network telescope data does appear to move in trends the concepts of supply and demand do not make sense. While there is a strong supply of network telescope data there is no demand for any for malicious activity.

| Metrics            |                |                 |                 |                  |            |
|--------------------|----------------|-----------------|-----------------|------------------|------------|
| Incident           | Price Channels | Bollinger Bands | Kelter Channels | Moving Envelopes |            |
| DDoS A             | Identified     | Identified      | Identified      | Identified       | Identified |
| DDoS B             | Identified     | Identified      | Identified      | Identified       | Identified |
| DDoS C             | Identified     | Failed          | Identified      | Identified       | Identified |
| DDoS D             | Identified     | Identified      | Identified      | Identified       | Identified |
| DDoS E             | Identified     | Failed          | Identified      | Identified       | Identified |
| Conficker A        | Identified     | Identified      | Identified      | Identified       | Identified |
| Conficker B        | Identified     | Failed          | Identified      | Identified       | Identified |
| Conficker C        | Identified     | Failed          | Identified      | Identified       | Identified |
| Conficker D        | Identified     | Failed          | Identified      | Identified       | Failed     |
| Conficker E        | Identified     | Failed          | Identified      | Identified       | Failed     |
| Net Send Malware A | Identified     | Identified      | Failed          | Identified       | Identified |
| Net Send Malware B | Identified     | Identified      | Failed          | Identified       | Identified |
| Net Send Malware C | Identified     | Failed          | Failed          | Identified       | Identified |
| Net Send Malware D | Identified     | Failed          | Failed          | Identified       | Identified |
| Scanning A         | Failed         | Identified      | Identified      | Failed           | Failed     |
| Scanning B         | Failed         | Identified      | Identified      | Failed           | Failed     |
| Scanning C         | Identified     | Identified      | Identified      | Failed           | Failed     |
| Scanning D         | Failed         | Identified      | Identified      | Failed           | Failed     |

Table 7.6: Identification Ability of Technical Analysis.

Further, many of the indicators used in technical analysis use the stock market values of close, high and low of a particular quantity. As close has no analogous concept in network telescope data, there is no need for the telescope to close at the end of the day like the stock exchange does.

Further, technical analysis assumes that a value will not suddenly drop to zero as is often the case for many of the metrics derived from network telescope data such as number of packets received for a particular port.

## 7.3 Recommendations

Having considered a selection of metrics and techniques this final section articulates the recommendation of the researchers with respect to the application of these metrics and techniques to network telescope datasets.

The total packet count metric is an important metric and should be included in any form of network telescope analysis, the weakness of lack of specificity can be remediated by combination of other more specific metrics such as datagram counts. The frequency of source IPs, destination IPs, packet sizes and TTLs would need to consider spaces which are simply too large. The average packet size metric has shown to be effective in identifying a number of incidents, particularly scanning type activity and as such would be an asset to a monitoring system. The average packets per source IP seems to be an effective measure and most useful for identifying DDoS backscatter and scanning.

The only baseline type that seems to make sense from the research conducted is the rolling baseline. Rolling baselines that use averages seem sensible though choosing a period is tricky. The six-month rolling average was shown in this thesis to be a useful baseline though further research is suggested.

Technical analysis techniques that rely on the close, high and low values such as Keltner channels simply do not work within the network telescope context as there are no analogous concepts in network telescope data. As such an technical analysis concepts that rely on these measures are most likely to be unusable in the network telescope context. Moving average envelopes due to their percentage based nature are inappropriate for dealing with activity that grows exponentially as is the case for most worms.

## 7.4 Summary

This chapter evaluated the approaches considered in Chapters 5 and 6 by considering how these metrics and identification tools reacted to cases of DDoS, scanning, net send malware and Conficker. Further, the weakness, strengths and suitability of these metrics and techniques for use in the analysis of network telescope data was discussed.

Section 7.1 discussed how effective the total packet count, counts per packet size, counts per specific UDP or TCP destination port and ICMP type, and counts by source and destination IP are at identifying incidents.

Section 7.2 looked at the use of baselines and technical analysis techniques as a way to assist in the identification of incidents from metrics. Baseline analysis considered one-year recurring, six-month rolling average and one-month rolling maximum baselines. Technical analysis considered price channels, Bollinger bands, Keltner channels and moving average envelopes.

Finally, Section 7.3 made recommendations on the use of the metrics and techniques discussed in this chapter for network telescope analysis. The metrics recommended included: total packet count, counts per datagram type, average packet size, average destination IP and OS groupings. Rolling six-month averages baselines were suggested to be used by all metrics to allow for easier inference of abnormal activity. Further research into technical analysis indicators formed from the high, low and close values of the stock market was not suggested unless a mapping between these values and values in the network telescope data could be formed. Bands such as Bollinger bands and price channels were suggested as techniques which would be useful in network telescope data though future research.

*I have an almost religious zeal - not for technology per se, but for the Internet which is for me, the nervous system of mother Earth, which I see as a living creature, linking up.*

Dan Millman

# 8

## Conclusion

**T**HIS thesis has studied approaches to analysing data obtained from a network telescope to aid in detecting potential network incidents. This was achieved by studying the properties and statistical nature of network telescope data and then designing a number of metrics and applying data analysis techniques from existing fields of study. These approaches and metrics were evaluated by considering how effective they were at detecting previously identified malicious activity.

Having completed the main body of work it is now necessary to conclude the findings made in this thesis and suggest how to further the field of network telescope research. This closing section provides a summary of the work conducted, reviews the research objectives, discusses the limitations of the work presented and finally suggests future work in this field.

### 8.1 Summary of Previous Chapters

A brief overview of the work conducted in this thesis is now presented.

**Chapter 1** defined the context, objectives and focus of the research conducted in this thesis.

**Chapter 2** provided the necessary background on matters related to network telescopes and data analysis. In particular the topics of network security, network telescopes and analysis techniques from existing fields were reviewed.

**Chapter 3** discusses the properties of the datasets studied and the significant variables associated with network telescope data. A simple descriptive statistical study was performed to aid in the selection of metrics and identification techniques.

**Chapter 4** examined the datasets for incidents for use in the evaluation phase of this thesis. Incidents identified included DDoS, Conficker, port scanning and netsend malware.

**Chapter 5** explored a number of metrics that can be formed from network telescope data by plotting the metric for varying time periods and observing the metrics behaviour. Metrics were formed from counts, aggregation and grouping.

**Chapter 6** considered a number of data analysis techniques by applying said techniques to the network telescope data and observing the result. A selection of techniques from baseline analysis and technical analysis were applied to the datasets and the behaviour of these identification techniques was observed.

**Chapter 7** evaluated the metrics and analysis techniques by applying them to the incidents outlines in Chapter 4. It was considered whether these approaches were capable of identifying this abnormal activity. The advantages, disadvantages and general opinion of the researchers towards certain approaches was provided as well.

## 8.2 Review of the Research Objectives

The research objectives from Chapter 1 are now re-evaluated by considering the work achieved.

- Through observation of the telescope data, as described in Chapter 3, an initial understanding of network telescope data was formed through descriptive statistics.
- Cursory hand parses of the packet captures provided by the telescopes yielded a number of incidents that could be used to test metrics and analytical approaches.

These incidents were documented in both time period and behaviour and may be useful to other researchers in need of such data. Incidents identified included DDOS, Conficker, port scanning and netsend malware.

- Metrics were derived from network telescope data and a number of them were shown to be useful in identifying anomalous activity. These metrics included:
  - Total packet counts.
  - Frequency of specific packet sizes.
  - Packets received per destination port or ICMP type.
  - Packets received per datagram type.
  - Average Packet Size.
  - Average TTL.
  - Average Counts per Source IP.
  - Average Counts per Destination IP.
  - Ratio of packets received per datagram type.
  - Packets received per OS Grouping.
  - Counts from bogon IPs.
- A number of new techniques from other fields were explored for network telescope analysis. Baseline analysis was considered and shown to be somewhat effective in identifying unusual activity with rolling averages chosen as the most successful of these techniques.

The field of technical analysis was considered with evaluations of moving average envelopes, Donchian price channels, Bollinger bands and Keltner channels. While it was shown that technical analysis does not form a perfect fit for network telescope analysis due to the mismatch of technical assumptions and the reality of data obtained from network telescope it appears that some of these techniques with modification may still be useful.

### 8.3 Limitations of Work

It is important to note that the work described in this thesis is not definitive and only provides a starting point from which a more complete study of network telescope data analysis could be formed. A number of these limitations are discussed below:

- The number of datasets considered is small due to the time intensive process of data analysis and limited access to data. A more comprehensive study of network telescope data and approaches is needed to advance the legitimacy of the field.
- Deciding whether malicious activity is worthy of being defined as an incident is particularly challenging as there is currently no agreed upon standard in any literature known to the researchers. Even if there was it would most likely be highly contested. This is further complicated by the fact that new strains of malicious activity may very well break any definitions created.
- It is difficult to assess how effective a technique is by considering the ability of a technique to identify malicious activity as defining malicious activity is challenging. This thesis has only considered a limited number of incidents and applying the techniques for a small number of time windows. It is all together possible that some of the less effective techniques may be useful at identifying other types of behaviour under different circumstances and vice versa.
- Traffic analysis and technical analysis share a common problem in that both fields deal with a considerable amount of uncertainty. For example, it may be possible that movements in the market are not indicative of a trend at all and much the same changes in composition of malicious traffic may not actually reflect the outbreak of a new worm or virus.
- Heavy periods of traffic in December 2006 illustrate a case where there is a clear increase in traffic, however research yield very little evidence of an attack occurring. This illustrates a weakness of using data from telescopes as a way of testing the effectiveness of metrics for identification of attacks. Further, this data is so complex, it is difficult to find a control group for comparison.

## 8.4 Future Work

As network telescope data analysis is a young field it is a mostly uncharted field of study with many potential avenues of research. The work presented in this thesis describe new approaches to network telescope monitoring and incident identification. While the research is rather broad in terms of areas of research which are considered there is still a large margin for improvement and other areas to consider application from. At the start of the research and design of the research goals and question a number of approaches were

considered but not examined. These excluded approaches together with items of interest that resulted from this thesis are now considered as potential future projects:

- A study of the effect that aggregating data collected by multiple network telescopes nodes, to allow for more distributed monitoring, has on the techniques described. It is possible that certain techniques may be more useful with more aggregated data while other techniques fail in such cases.
- There appears to be a disjoint between the assumptions of technical analysis and what holds true for network telescope data resulting in some unusual reactions from certain indicators. It may be interesting to derive these indicators under different assumptions which hold true for network telescope data.
- 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.
- Some of the techniques described, especially those from technical analysis, may be appropriate for identification of events in other monitor systems such as internal monitoring of an organization network structure to identify infected hosts.
- Only a limited number of data analysis techniques were considered in this thesis. Other interesting fields of study that may be of use to network telescope analysis include spectral analysis,

In closing, it is important to note that there are a number of challenges still remaining with the use of network telescopes as a means of identifying malicious activity on the Internet. While some of these issues are irresolvable, malicious activity cannot be inferred from a network telescope there is little to be gained from analysing telescope data, it is important to note that network telescope should be used in conjunction with other monitoring methods to create a more comprehensive approach to monitoring malicious activity on the Internet. The core issue of complexity and data analysis has been addressed in this thesis by considering ways to construct metrics and identification techniques however the researchers acknowledge that approaches considered in this thesis need a great deal of further study and refinement.

## References

- Aben E.** *Conficker/Conflicker/DownAdUp as Seen From the UCSD Network Telescope*. [On-line]. Available: <http://www.caida.org/research/security/ms08-067/conficker.xml>. 27 February 2009. [Last accessed: 4 August 2010].
- Alexander B.** *Intrusion Detection FAQ: Port 137 Scan*. [On-line]. Available: [http://www.sans.org/security-resources/idfaq/port\\_137.php](http://www.sans.org/security-resources/idfaq/port_137.php). 10 May 2000. [Last accessed: 17 August 2011].
- Anderson N.** *Source: Anonymous Attacks on Sony Annoying, Not Much More*. [On-line]. Available: <http://arstechnica.com/tech-policy/news/2011/04/source-anonymous-attacks-on-sony-annoying-not-much-more.ars>. 10 April 2010. [Last accessed: 4 April 2011].
- Baker F., Harrop W., and Armitage G.** *IPv4 and IPv6 Greynets*. [On-line]. Available: <http://www.ietf.org/rfc/rfc6018.txt>. September 2010. [Last accessed: 5 October 2011].
- British Broadcasting Corporation (BBC).** *Police Investigate Habbo Hotel Virtual Furniture Theft*. [On-line]. Available: <http://www.bbc.co.uk/news/10207486>. 1 June 2010. [Last accessed: 28 July 2011].
- Beck F., Festor O., and State R.** *High Security Laboratory - Network Telescope Technical Report* [On-line]. Available: [http://hal.archives-ouvertes.fr/docs/00/33/75/68/PDF/Technical\\_Report\\_Network\\_Telescope.pdf](http://hal.archives-ouvertes.fr/docs/00/33/75/68/PDF/Technical_Report_Network_Telescope.pdf). March 2007. [Last accessed: 8 October 2011].
- Bethencourt J., Low J., Simmons I. and Williamson M** “Establishing Darknet Connections: An Evaluation of Usability and Security”, *In the Proceedings of the 3<sup>rd</sup> Symposium on Usable Privacy and Security*. 2007. [Last accessed: 14 September 2010]. <http://doi.acm.org/10.1145/1280680.1280700>.

- Bidgoli H.** *Handbook of Information Security*. USA, California: Wiley Press. 1<sup>st</sup> ed., vol 3. 2006. ISBN: 978-047164833-8.
- Bollinger J.** *Bollinger on Bollinger Bands*. USA: McGraw-Hill. 1<sup>st</sup> ed. 2002a. ISBN: 978-0071373685.
- Bollinger J.** “Construction” in *Bollinger on Bollinger Bands*. USA: McGraw-Hill. 1<sup>st</sup> ed. 2002b. ISBN: 978-0071373685.
- Bollinger J.** “Bollinger Band Indicators” in *Bollinger on Bollinger Bands*. USA: McGraw-Hill. 1<sup>st</sup> ed. pp 60–67. 2002c. ISBN: 978-0071373685.
- Co-operative Association for Internet Data Analysis (CAIDA).** *Passive Data Collection: UCSD Network Telescope*. [On-line]. Available: [http://www.caida.org/data/passive/network\\_telescope.xml](http://www.caida.org/data/passive/network_telescope.xml). January 2010a. [Last accessed: 19 April 2010].
- Co-operative Association for Internet Data Analysis (CAIDA).** *CAIDA’s Chicago A Passive Network Monitor, Application (packets/second)*. [On-line]. Available: [http://www.caida.org/data/realtime/passive/?monitor=equinix-chicago-dirA&row=timescales&col=sources&sources=app&graphs\\_sing=ts&counters\\_sing=packets&timescales=24&timescales=168&timescales=672&timescales=17520](http://www.caida.org/data/realtime/passive/?monitor=equinix-chicago-dirA&row=timescales&col=sources&sources=app&graphs_sing=ts&counters_sing=packets&timescales=24&timescales=168&timescales=672&timescales=17520). January 2010b. [Last accessed: 19 April 2010].
- Computer Emergency Response Team (CERT).** *Exploitation of Unprotected Windows Networking Shares*. [On-line]. Available: [http://www.cert.org/incident\\_notes/IN-2000-02.html](http://www.cert.org/incident_notes/IN-2000-02.html). 7 April 2000. [Last accessed: 12 September 2011].
- Computer Emergency Response Team (CERT).** *Advisory CA-2003-04 MS-SQL Server Worm*. [On-line]. Available: <http://www.cert.org/advisories/CA-2003-04.html>. 27 January 2003. [Last accessed: 4 September 2011].
- Chart School.** Price Channels [On-line]. Available: [http://stockcharts.com/help/doku.php?id=chart\\_school:technical\\_indicators:price\\_channels](http://stockcharts.com/help/doku.php?id=chart_school:technical_indicators:price_channels). 2010. [Last accessed: 3 April 2010].
- Claffy K., Hyun Y., Keys K., Fomenkov M. and Krioukov D.** “Internet Mapping: From Art to Science”. In *the Proceedings of the Cybersecurity Applications & Technology Conference for Homeland Security 2009*. pp 47–59. <http://dx.doi.org/10.1109/CATCH.2009.38>. 2009.

- Colby R.** *The Encyclopaedia Of Technical Market Indicators*. USA: McGraw-Hill. 2<sup>nd</sup> ed. vol 1. 2002. ISBN: 978-1556230493.
- Conficker Working Group.** *Conficker Timeline*. [On-line]. Available: <http://www.confickerworkinggroup.org/wiki/pmwiki.php/ANY/Timeline>. 26 April 2009. [Last accessed: 3 April 2011].
- Corrons L.** *4chan Users Organize Surgical Strike Against MPAA*. [On-line]. Available: <http://pandalabs.pandasecurity.com/4chan-users-organize-ddos-against-mpaa/>. 17 August 2010. [Last accessed: 4 April 2011].
- Coursey W.** “Descriptive Statistics: Summary Numbers” in *Statistics and Probability for Engineering Applications*. Canada: University of Saskatchewan Press. pp 41-51. 2003. ISBN: 978-0750676182.
- Cowie B. and Irwin B.** “Data Classification for Artificial Intelligence Construct Training to Aid in Network Incident Identification Using Network Telescope Data”. In *the Proceedings of the Annual Research Conference of the South African Institute of Computer Scientists and Information Technologists (SAICSIT) 2010*. Bela Bela, South Africa. October 2010a. <http://dx.doi.org/10.1145/1899503.1899544>.
- Cowie B. and Irwin B.** “A Baseline Numeric Analysis of Network Telescope Data for Network Incident Discovery”. In *the Proceedings of the 13<sup>th</sup> Southern African Telecommunication and Applications Conference (SATNAC)*. Spier Wine Estate, Stellenbosch, South Africa. October 2010b. <http://www.satnac.org.za/proceedings/2010/papers/management/Cowie%20FP%20v3%20442.pdf>.
- Cowie B. and Irwin B.** “An Evaluation of Trading Bands as Indicators for Network Telescope Datasets”. In *the Proceedings of the 14<sup>th</sup> Southern African Telecommunication and Applications Conference (SATNAC)*. East London International Convention Centre, East London, South Africa. September 2011c.
- Danchev D.** *Conficker’s Estimated Economic Cost? \$9.1 billion*. [On-line]. Available: <http://www.zdnet.com/blog/security/confickers-estimated-economic-cost-91-billion/3207>. 23 April 2009. [Last accessed: 15 June 2011].
- Dell Computer Corporation.** *Why Analyze Firewall Logs?* [On-line]. Available: [http://www.secureworks.com/research/articles/other\\_articles/firewall-primer/](http://www.secureworks.com/research/articles/other_articles/firewall-primer/). 23 April 2009. [Last accessed: 15 June 2011].

- Eddy W.** *TCP SYN Flooding Attacks and Common Mitigations*. [On-line]. Available: <http://tools.ietf.org/html/rfc4987/>. August 2007. [Last accessed: 11 September 2011].
- Edwards R. and Magee J.** *Technical Analysis of Stock Trends*. USA, Florida: CRC Press. 8<sup>th</sup> ed. 2001. ISBN: 978-0814406809.
- Faires J. and Burden R.** *Numerical Analysis*. USA: Thompson. 8<sup>th</sup> ed. 2005. ISBN: 978-0538733519.
- Fuchs C.** *Internet and Society: Social Theory in the Information Age*. Routledge. 1<sup>st</sup> ed. 2008. ISBN: 978-0415961325.
- Gruener W.** *Windows Market Share Drops to 15-year Low*. [On-line]. Available: <http://www.tgdaily.com/trendwatch-features/40398-windows-market-share-drops-to-15-year-low>. 2008. [Last accessed: 11 January 2012].
- Harder U., Johnson M., Bradley J. and Knottenbelt W.** “Observing Internet Worm and Virus Attacks with a Small Network Telescope”. *Electronic Notes in Theoretical Computer Science*. vol 151(3). pp 47–59. 2004. <http://dx.doi.org/10.1016/j.entcs.2006.03.011>.
- Harrop W. and Armitage G.** “Defining and Evaluating Greynets (Sparse Darknets)”. *In the Proceedings of the IEEE Conference on Local Computer Networks 30<sup>th</sup> Anniversary*. pp 344–350. 2005. <http://dx.doi.org/10.1109/LCN.2005.46>.
- Holman V., Johnson M. and Bradley J.** *Creating Performance Baselines That Establish Goals and Standards*. [On-line]. Available: <http://www.slideshare.net/victorholman/creating-performance-baselines-that-establish-goals-and-standards-presentation>. 2009. [20 April 2010].
- Hyun Y.** “Archipelago Measurement Infrastructure: Status and Experiences”. *In the Proceedings of the 10<sup>th</sup> CAIDA-WIDE-CASFI Workshop*. August 2008. [http://www.caida.org/publications/presentations/2008/wide\\_young\\_ark/wide\\_young\\_ark.pdf](http://www.caida.org/publications/presentations/2008/wide_young_ark/wide_young_ark.pdf).
- Internet Assigned Numbers Authority (IANA).** *IANA IPv4 Address Space Registry*. [On-line]. Available: <http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml>. 3 February 2011. [Last accessed: 1 December 2010].

- Internet Crime Compliant Center (IC<sup>3</sup>).** *Internet Crime Report - 2009*. [On-line]. Available: [http://www.ic3.gov/media/annualreport/2009\\_IC3Report.pdf](http://www.ic3.gov/media/annualreport/2009_IC3Report.pdf). 2009. [Last accessed: 21 April 2010].
- Information Sciences Institute.** *Internet Protocol: Darpa Internet Program Protocol Specification*. RFC 791. [On-line]. Available: <http://www.ietf.org/rfc/rfc791.txt>. September 1981a. [Last accessed: 14 April 2010].
- Information Sciences Institute.** *Transmission Control Protocol*. RFC 793. [On-line]. Available: <http://www.ietf.org/rfc/rfc791.txt>. September 1981b. [Last accessed: 14 April 2010].
- Internet World Stats.** *Internet Growth Statistics*. [On-line]. Available: <http://www.internetworldstats.com/stats.htm>. 2011. [Last accessed: 14 April 2011].
- Irwin B.** “Conficker: 687 days later”. *Presentation at the 2<sup>nd</sup> Annual ZaCon Security Conference*. University of Johannesburg, South Africa. 2010. [http://www.zacon.org.za/Archives/2010/slides/2010\\_ZaCon\\_Barry\\_Irwin.pdf](http://www.zacon.org.za/Archives/2010/slides/2010_ZaCon_Barry_Irwin.pdf).
- Irwin B.** “A Framework for the Application of Network Telescope Sensors in a Global IP Network”. PHD thesis. Rhodes University, Grahamstown, South Africa. 2011.
- Jaquith A.** *Security Metrics*. USA: Addison-Wesley. 1<sup>st</sup> ed. 2007. ISBN: 978-0321349989.
- Keys K., Moore D., Koga R., Lagache E. and Claffy K.** “The Architecture of CoralReef: An Internet Traffic Monitoring Software Suite”. *In the Proceedings of the Passive & Active Measurement Workshop*. 2001. <http://www.caida.org/publications/papers/2001/CoralArch/coralreef.pdf>.
- Kirkpatrick C. and Dalhquist J.** “Moving Averages” in *Technical Analysis: The Complete Resource for Financial Market Technicians*. Pearson Education. 2<sup>nd</sup> ed. 2011. ISBN: : 978-0137059447.
- Kitchens L.** “Types of Data” in *Exploring Statistics*. USA: Thomson Publishing Company. 2<sup>nd</sup> ed. 2002. ISBN: 978-0314284983.
- Kumar R. and Kush A.** “E-Learning Emergence”. *DESIDOC Bulletin of Information Technology*. vol 26. pp 19–24. March 2006. <http://publications.drdo.gov.in/ojs/index.php/djlit/article/download/93/21>.
- Leiner B., Cerf V., Clark D., Kahn R., Kleinrock L., Lynch D., Postel J., Roberts L. and Wolff S.** *A Brief History of the Internet*. [On-line]. Available: <http://www.ietf.org/rfc/rfc791.txt>.

- [//www.isoc.org/internet/history/brief.shtml](http://www.isoc.org/internet/history/brief.shtml). 2002. [Last accessed: 20 October 2011].
- Lemos R.** *Counting the Cost of Slammer*. [On-line]. Available: <http://news.cnet.com/2100-1001-982955.html>. 31 Jan 2003. [Last accessed: 1 April 2011].
- Lewine D.** *POSIX Programmers Guide*. USA: O'Reilly Media. 1<sup>st</sup> ed. 1991. ISBN: 978-0937175736.
- Litchfield D.** *The Inside Story of SQL Slammer*. [On-line]. Available: [http://threatpost.com/en\\_us/blogs/inside-story-sql-slammer-102010](http://threatpost.com/en_us/blogs/inside-story-sql-slammer-102010). 20 October 2010. [Last accessed: 11 September 2011].
- Maurer J.** *Internet Worms: Walking on Unstable Ground*. [On-line]. Available: [http://www.sans.org/reading\\_room/whitepapers/malicious/internet-worms-walking-unstable-ground\\_1229](http://www.sans.org/reading_room/whitepapers/malicious/internet-worms-walking-unstable-ground_1229). 2003. [Last accessed: 1 April 2011].
- Mick J.** *Anonymous Engages in Sony DDoS Attacks Over GeoHot PS3 Lawsuit*. [On-line]. Available: <http://www.dailytech.com/Anonymous+Engages+in+Sony+DDoS+Attacks+Over+GeoHot+PS3+Lawsuit/article21282.htm>. 4 April 2009. [Last accessed: 11 September 2011].
- Microsoft.** *Windows XP Documentation - Net send*. [On-line]. Available: [http://www.microsoft.com/resources/documentation/windows/xp/all/proddocs/en-us/net\\_send.mspx?mfr=true](http://www.microsoft.com/resources/documentation/windows/xp/all/proddocs/en-us/net_send.mspx?mfr=true). 2001. [Last accessed: 15 September 2011].
- Microsoft.** *Messenger Service Window that Contains an Internet Advertisement Appears*. [On-line]. Available: <http://support.microsoft.com/kb/330904>. 26 January 2011. [Last accessed: 11 September 2011].
- Mircosoft.** *Microsoft Security Bulletin MS08-067 - Critical*. [On-line]. Available: <http://technet.microsoft.com/en-us/security/bulletin/ms08-067>. 23 October 2008. [Last accessed: 11 September 2011].
- Nison S.** *Core Concepts in Data Analysis: Summarization, Correlation and Visualization*. Springer. 1<sup>st</sup> ed. 2011. ISBN: 978-0857292865.
- Mokube I. and Adams M.** "Honeypots: Concepts, Approaches, and Challenges", *Proceedings of the 45<sup>th</sup> Annual Southeast Regional Conference*. pp 321–326. 2007. <http://dx.doi.org/10.1145/1233341.1233399>.

- Moore D.** *Network Telescopes: Observing Small or Distant Security Events*. [On-line]. Available: [http://www.caida.org/publications/presentations/2002/usenix\\_sec/](http://www.caida.org/publications/presentations/2002/usenix_sec/). 8 August 2002. [Last accessed: 25 August 2010].
- Moore D., Keys K., Koga R., Lagache E. and Claffy K.** “The CoralReef Software Suite as a Tool for System and Network Administrators”. In *the Proceedings of the 15<sup>th</sup> USENIX Conference on System Administration*. pp 133–144. 2001. <http://www.caida.org/outreach/papers/2001/CoralApps/CoralApps.pdf>.
- Moore D., Paxson V., Savage S., Shannon C., Staniford S. and Weaver, N.** “Inside the Slammer Worm”. *IEEE Security and Privacy*. vol 1(4). pp 33–39. 2003. <http://dx.doi.org/10.1109/MSECP.2003.1219056>.
- Moore D. and Shannon C.** “The Spread of the Witty Worm”. *IEEE Security and Privacy*. vol 2. pp 46–50. July 2004. <http://dl.acm.org/citation.cfm?id=1018027.1018275>.
- Moore D., Shannon C. and Claffy K.** “Code-Red: A Case Study on the Spread and Victims of an Internet Worm”. In *the Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement*. pp 273–284. 2002. <http://dx.doi.org/10.1145/637201.637244>.
- Moore D., Shannon C., Voelker G. and Savage S.** *Network Telescopes: Technical Report*. [On-line]. Available: <http://www.caida.org/publications/papers/2004/tr-2004-04/tr-2004-04.pdf>. [Last accessed: 10 April 2010].
- Moore D., Voelker G. and Savage S.** “Inferring Internet Denial-Of-Service Activity”. *ACM Transactions on Computer Systems*. vol 24(2). 2006. <http://dx.doi.org/10.1145/1132026.1132027>.
- Murphy J.** *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*. New York City, New York: New York Institute of Finance. 1999a. ISBN: 978-0735200661.
- Murphy J.** “Basic Concepts of Trend” in *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*. New York City, New York : New York Institute of Finance. pp 55–65. 1999b. ISBN: 978-0735200661.
- Nahorney B.** *The DownAdUp Codex*. [On-line]. Available: [http://www.symantec.com/content/en/us/enterprise/media/security\\_response/whitepapers/the\\_downadup\\_codex\\_ed2.pdf](http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/the_downadup_codex_ed2.pdf). 2 June 2009. [Last accessed: 8 May 2011].

- National Institute of Standards and Technology (NIST).** *Vulnerability Summary for CVE-2008-1701*. [On-line]. Available: <http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2008-1701>. 8 August 2008. [Last accessed: 5 May 2011].
- Organisation for Economic Co-operation and Development (OECD).** *The E-government Imperative: Main Findings*. [On-line]. Available: <http://www.oecd.org/dataoecd/60/60/2502539.pdf>. March 2003. [Last accessed: 10 June 2011].
- Orman H.** “The Morris Worm: A Fifteen-Year Perspective”. *IEEE Security and Privacy*. vol 1(5). pp 35–43. September 2003. <http://dx.doi.org/10.1109/MSECP.2003.1236233>.
- Pang R., Yegneswaran V., Barford P., Paxson V. and Peterson L.** “Characteristics of Internet Background Radiation”, *In the Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement*. pp 27–40. 2004. <http://dx.doi.org/10.1145/1028788.1028794>.
- Payne S.** *A Guide to Security Metrics*. [On-line]. Available: [http://www.sans.org/reading\\_room/whitepapers/auditing/guide-security-metrics\\_55/](http://www.sans.org/reading_room/whitepapers/auditing/guide-security-metrics_55/). 19 June 2006. [Last accessed: 7 May 2011].
- Pemberton D., Komisarczuk P. and Welch I.** “Internet Background Radiation Arrival Density and Network Telescope Sampling Strategies”, *In the Proceedings of the Telecommunication Networks and Applications Conference*. pp 246–252. December 2007. <http://dx.doi.org/10.1109/ATNAC.2007.4665254>.
- Porras P. and Yegneswaran V.** “A Foray into Confickers Logic and Rendezvous Points”. *In the Proceedings of the 2<sup>nd</sup> USENIX Workshop on Large-Scale Exploits and Emergent Threats*. vol 2. pp 7. 2009. <http://dl.acm.org/citation.cfm?id=1855676.1855683>.
- Porras P., Saidi H. and Yegneswaran V.** *An Analysis of Conficker’s Logic and Rendezvous Points*. [On-line]. Available: <http://mtc.sri.com/Conficker/>. 19 March 2009. [Last accessed: 15 June 2011].
- J Postel.** *User Datagram Protocol*. RFC 768 . [On-line]. Available: <http://www.ietf.org/rfc/rfc768.txt>. August 1980. [Last accessed: 14 April 2010].
- J Postel.** *Internet Control Message Protocol: Darpa Internet Program Protocol Specification*. RFC 792. [On-line]. Available: <http://www.ietf.org/rfc/rfc792.txt>. September 1981. [Last accessed: 14 April 2010].

- Preece J., Maloney-Krichmar D. and Abras C.** “History of Emergence of Online Communities”. In *Encyclopedia of Community* . 2003. <http://www.ifsm.umbc.edu/preece/paper/6%20Final%20Enc%20preece%20et%20a1.pdf>.
- Press W., Teukolsky S., Vetterling W. and Flannery B.** “Filtering” in *Numerical Recipes in C, The Art of Scientific Computing, Second Edition*. UK: Cambridge University Press. pp 559–560. 1992. ISBN: 978-0521431088.
- Qiming L., Xu R. and Wang S.** “Modelling and Analysis of an SIRS Model for Worm Propagation”. In *the Proceedings of Computational Intelligence and Security*. pp 361–365. December 2009. <http://dx.doi.org/10.1109/CIS.2009.187>.
- Ray E.** *Malware FAQ: MS-SQL Slammer*. [On-line]. Available: <http://www.sans.org/security-resources/malwarefaq/ms-sql-exploit.php>. 2004. [Last accessed: 7 May 2011].
- Reuters.** *Hackers hit Hollywood’s piracy watchdog*. [On-line]. Available: <http://www.reuters.com/article/2010/09/20/us-hackers-idUSTRE68J09F20100920>. 19 September 2010. [Last accessed: 21 May 2011].
- Scacchi W.** *The Emergence of Electronic Commerce on the Internet*. USA: USC Business. vol 5. pp 32–34. 1994.
- Scheifler R.** *RFC 1013 - X Window System Protocol*. [On-line]. Available: <http://tools.ietf.org/html/rfc1013>. June 1987. [Last accessed: 14 September 2011].
- Schneier B.** *Secrets & Lies - Digital Security in a Networked World*. USA, New York: John Wiley & Sons, Inc. 1<sup>st</sup> ed. 2000. ISBN: 978-0471253112.
- Shinoda Y., Ikai K. and Itoh M.** “Vulnerabilities of Passive Internet Threat Monitors”. In *the proceedings of the 14<sup>th</sup> conference on USENIX Security Symposium*. 2005. <http://dl.acm.org/citation.cfm?id=1251398.1251412>.
- Team Cymru.** *The Bogon Reference*. [On-line]. Available: <http://www.team-cymru.org/Services/Bogons/>. 15 August 2011a. [20 August 2011].
- Team Cymru.** *Team Cymru - About*. [On-line]. Available: <http://www.team-cymru.org/About/>. 15 August 2011b. [20 August 2011].
- Tsotsis A.** *RIAA Goes Offline, Joins MPAA as Latest Victim of Successful DDoS Attacks*. [On-line]. Available: <http://techcrunch.com/2010/09/19/riaa-attack/>. 19 September 2010. [20 August 2011].

- US Computer Emergency Response Team (US-CERT).** *Conficker Worm Targets Microsoft Windows Systems*. [On-line]. pp 1. Available: <http://www.us-cert.gov/cas/techalerts/TA09-088A.html>. 29 March 2009a. [Last accessed: 5 May 2011].
- US Computer Emergency Response Team (US-CERT).** *Understanding Denial-of-Service Attacks*. [On-line]. Available: <http://www.us-cert.gov/cas/tips/ST04-015.html>. 4 November 2009b. [Last accessed: 11 September 2011].
- van Riel J.-P. and Irwin B.** “InetVis, A Visual Tool for Network Telescope Traffic Analysis”. In *the Proceedings of the 4<sup>th</sup> International Conference on Computer Graphics, Virtual Reality, Visualization and Interaction in Africa (AFRIGRAPH)*. pp 485–489. 2006a. <http://dx.doi.org/10.1145/1108590.1108604>.
- van Riel J.-P. and Irwin B.** “Identifying and Investigating Intrusive Scanning Patterns by Visualizing Network Telescope Traffic in a 3-D Scatter-Plot”. In *the Proceedings of the 6<sup>th</sup> Annual Information Security South Africa (ISSA) Conference*. Balalaika Hotel, Sandton, South Africa. July 2006b.
- Vijayan J.** *Oak Ridge National Lab Shuts Down Internet, Email after Cyberattack*. [On-line]. Available: <http://www.networkworld.com/news/2011/041911-oak-ridge-national-lab-shuts.html>. 15 September 2010. [20 August 2011].
- Wagenseil P.** *Anonymous Hacktivists Attack Egyptian Websites*. [On-line]. Available: [http://www.msnbc.msn.com/id/41280813/ns/technology\\_and\\_science-security/t/anonymous-hacktivists-attack-egyptian-websites/](http://www.msnbc.msn.com/id/41280813/ns/technology_and_science-security/t/anonymous-hacktivists-attack-egyptian-websites/). 26 January 2011. [Last accessed: 11 September 2011].
- Weaver R.** “A Probabilistic Population Study of the Conficker-C Botnet”. In *the Proceedings of the 11<sup>th</sup> International Conference on Passive and Active Measurement*. pp 181–190. 2010a. <http://www.cert.org/netsa/publications/weaver-conficker.pdf>.
- Wustrow E., Karir M., Bailey M., Jahanian F. and Huston G.** “Internet background radiation revisited”. In *the Proceedings of the 10<sup>th</sup> Annual Conference on Internet Measurement*. pp 62–74. 2010b. <http://doi.acm.org/10.1145/1879141.187914>.
- Wilder J.** *New Concepts in Technical Trading Systems*. North Carolina, Greensboro: Trend Research. 1<sup>st</sup> ed. 1978. ISBN: 978-0894590276.
- Wireshark Forums.** *Unexplained NETBIOS Traffic*. [On-line]. Available: <http://ask.wireshark.org/questions/2824/unexplained-netbios-traffic>. 15 March 2011. [Last accessed: 11 September 2011].

- W3Schools.** *OS Platform Statistics*. [On-line]. Available: [http://www.w3schools.com/browsers/browsers\\_os.asp](http://www.w3schools.com/browsers/browsers_os.asp). 2011. [Last accessed: 11 January 2012].
- Ylonen T.** *RFC 4253 - The Secure Shell (SSH) Transport Layer Protocol*. [On-line]. Available: <http://tools.ietf.org/html/rfc4253>. January 2006. [Last accessed: 22 October 2011].
- Zetter K.** “How Digital Detectives Deciphered Stuxnet, the Most Menacing Malware in History”. [On-line]. Available: <http://www.wired.com/threatlevel/2011/07/how-digital-detectives-deciphered-stuxnet/all/1>. 11 July 2011. [Last accessed: 28 July 2011].
- Zou C., Gong W., Towsley D. and Gao L.** “The Monitoring and Early Detection of Internet Worms”. *IEEE/ACM Transactions on Networking*. vol 13(5). pp 961–974. October 2005. <http://dx.doi.org/10.1109/TNET.2005.857113>.

# Appendices



## Glossary

### Applications

***CoralReef*** - Framework designed to perform analysis on network telescope datasets.

***p0f*** - Application that provides OS identification from packets.

***NMAP*** - Network exploration tool primarily used for scanning.

***SQL*** - Structured Query Language.

***Wireshark*** - Packet capture and analysis program.

## Organizations

*AfriNIC* - The Registry of Internet Number Resources for Africa.

*Arbor* - Network monitoring and security solutions company.

*ARPANET* - Advanced Research Projects Agency Network.

*CAIDA* - Cooperative Association for Internet Data Analysis.

*IANA* - Internet Assigned Numbers Authority

*PREDICT* - Protected Repository for the Defence of Infrastructure Against Cyber Threats.

*Team Cymru* - Cooperative Association for Internet Data Analysis.

*USCD* - University of California, San Diego.

*US-CERT* - United States Computer Emergency Response Team.

## Networking and Related Concepts

*ACK* - Packets with TCP acknowledgement flag set.

*AS* - Autonomous Systems.

*Backscatter* - Traffic that is inferred from traffic from spoofed IP.

*CIDR* - Classless Inter-Domain Routing.

*CSRF* - Cross-Site Request Forgery.

***Darknets, Black Hole and Internet Sink*** - Synonyms for a passive network telescope.

***DDoS*** - Distributed Denial of Service.

***FIN*** - Packets with TCP finish flag set.

***GMT*** - Greenwich Mean Time.

***ICMP*** - Internet Control Message Protocol.

***LOIC*** - Low Orbit Ion Canon, tool used to perform denial of service.

***MTU*** - Maximum Transmission Unit.

***Passive Network Telescopes*** - A network telescope that does not acknowledge or respond to requests.

***POSIX*** - Portable Operating System Interface.

***RFC*** - Request for Comment.

***RPC*** - Remote Procedure Call.

***RST*** - Packets with TCP reset flag set.

***Smurfing*** - Packets with TCP Synchronize flag set.

***SSL*** - Secure Socket Layer.

***Spoofing*** - A technique employed by hackers that allows for host impersonation.

***SYN*** - Packets with TCP synchronize flag set.

***TCP*** - Transmission Control Protocol.

***TTL*** - Time To Live.

**UDP** - User Datagram Protocol.

**URG** - Packets with TCP urgent flag set.

**Zombies** - Hosts compromised by vulnerabilities by hackers often used to perform DDoS.

**XSS** - Cross-site scripting.

## Technical Analysis

**BB** - Bollinger Bandwidth.

**%b** - Bollinger Percentage Band.

**EMA** - Exponential Moving Average.

**MA** - Moving Average.

**MACD** - Moving Average Convergence Divergence.

**RSI** - Relative Strength Index.

**SMA** - Simple Moving Average.

# B

## Networking Concepts

### **B.1 Bogon IP List**

Bogon IPs are those that should not be routable as they have been allocated for other purposes. Table B.1 lists these bogon IP blocks as of the 18<sup>th</sup> of January 2012 (IANA, 2011; Team Cymru, 2011a).

### **B.2 Common Ports and ICMP Types**

Tables B.2 and B.3 list some of the commonly used TCP and UDP ports. As there are 130672 ports between the protocols combined a shortened list is provided. Table B.4 provides some information on the ICMP types used in this thesis.

| <b>Netblock</b> | <b>Usage</b>                              |
|-----------------|---|
| 0.0.0.0/8       | Reserved for Self-identification of hosts |
| 127.0.0.0/8     | Reserved for Local Loop-back              |
| 169.254.0.0/16  | Reserved for Link Local                   |
| 198.51.100.0/24 | Reserved for TEST-NET-2                   |
| 172.16.0.0/12   | Reserved for Private-Use Networks         |
| 192.0.2.0/24    | Reserved for TEST-NET-1                   |
| 192.88.99.0/24  | Reserved for 6to4 Relay Anycast           |
| 192.168.0.0/16  | Reserved for Private-Use Networks         |
| 203.0.113.0/24  | Reserved for TEST-NET-3                   |
| 224/8 - 255/8   | Reserved for Multicast                    |

Table B.1: Bogon IP Netblocks

| <b>Port Number</b> | <b>Usage</b>            |
|--------------------|-------------------------|
| 22                 | SSH                     |
| 80                 | HTTP                    |
| 139                | NetBIOS                 |
| 443                | SSL                     |
| 445                | Windows Network Sharing |
| 1433               | MSSQL Server            |
| 2967               | Symantec System Center  |
| 49152 - 65535      | Dynamic Port Range      |

Table B.2: Commonly Used TCP Ports

| <b>Netblock</b> | <b>Usage</b>        |
|-----------------|---------------------|
| 53              | DNS                 |
| 135             | RPC                 |
| 137             | NetBIOS             |
| 1026 - 1027     | Microsoft Messenger |
| 1434            | MSSQL Server        |

Table B.3: Commonly Used UDP Ports

| ICMP Type | Message                                     |
|-----------|---|
| 0         | Echo Reply                                  |
| 3         | Network error (host unknown or unreachable) |
| 5         | Redirect                                    |
| 8         | Echo Request                                |
| 11        | TTL Expired                                 |
| 12        | IP header error                             |
| 17        | Address mask request or reply               |
| 42 - 255  | Reserved                                    |

Table B.4: Commonly Used ICMP Types

## B.3 Operating Systems Identified from p0f

A list of the Operating Systems that can uniquely identified by p0f is provided below:

- Windows CE
- Windows NT 4.0 SP6a
- Windows 95 (b)
- Windows 98
- Windows 2000 with SP1 - P4
- Windows 2003 (AS)
- Windows XP with SP2/SP3
- Linux 2.0.3 - 2.6
- BSD 3.1 - 4.3
- FreeBSD 2-6
- OpenBSD 3.0 - 3.9
- Solaris 7 - 10
- MacOS 9.0 - 9.2
- Redline T—X 2200

- HP-UX B.10.20
- PocketPC 2002
- NetBSD
- Proxyblocker
- NMAP scans
- Eagle Secure Gateway
- Cisco 12008