# A System for Characterising Internet Background Radiation

**Literature Review**

David Yates

May 30, 2014

## 1 Introduction

This literature review introduces the reader to Internet Background Radiation (IBR), both its components and common recording practices. In doing so, it will set up a theoretical basis for the construction of the IBR characterisation system to follow.

The literature review is split into three main parts. Section 2 provides background information on IBR and its components: worms, scanning activities, backscatter from reflected distributed denial-of-service (DDoS) attacks, and misconfigurations. These are the sources of the data being characterised.

Section 3 discusses the Border Gateway Protocol (BGP), the protocol used on the internet for routing packets between Autonomous Systems (ASs). This is of interest as the construction of live BGP tables for the collected data may provide a way of discovering spoofed IP addresses.

Section 4 provides information on IBR collection tools and analysis methods. The collection tools discussed are darknets (also called network telescopes, sinkholes, blackhole monitors or background radiation monitors (Bailey *et al.*, 2006)) and greynets (Harrop & Armitage, 2005). In this literature review, darknets and greynets will both referred to as network telescopes, with the terms "darknet" or "greynet" being used when more specificity is required.

The analysis tools discussed are packet-level analysis, network flows and and honeynets. This section also delves into previous work in characterising network activities with statistical tools. This will form a basis for the system implemented in this study.

In this literature review, the CIDR notations /8, /16 and /24 will be used to refer to the traditional pre-CIDR Class A, Class B and Class C subnets (Fuller *et al.*, 1993) respectively, in order to maintain consistency with the referencing of differently sized IPv4 subnets (such as /19s). The number following the "/" denotes the number of fixed bits (out of 32 total bits) in each of the addresses within the subnet.

## 2 Internet Background Radiation

The term Internet Background Radiation (IBR) was coined by Pang *et al.* (2004) to describe an ongoing variety of unproductive network traffic destined for addresses not set up to receive it.

Traffic of this nature results from four main sources: worm and virus activity, network reconnaissance scans, backscatter from Distributed Denial of Service (DDoS) attacks and misconfigured networking equipment such as routers and servers (Pang *et al.*, 2004).

Pang *et al.* (2004) discovered that packets with the TCP SYN-ACK and TCP RST flags set made up the majority of the darknet data recorded over four days on /19 network, over one week on ten adjacent /24 networks and a over one week on a /8 network. This dominance can be seen in Figure 1, taken from Pang *et al.* (2004), which shows the data for the /8 network used.

Apart from that, the data recorded on each network had few similarities. The extreme volatility of the IBR in comparison to ordinary, productive traffic (that is, traffic made up by legitimate connections without malicious intent or misconfiguration on either side) was noted, as was the potential difficulty of discovering new types of traffic, especially new types of worms, as they may be intentionally designed to use the same ports as other worms (Pang *et al.*, 2004).

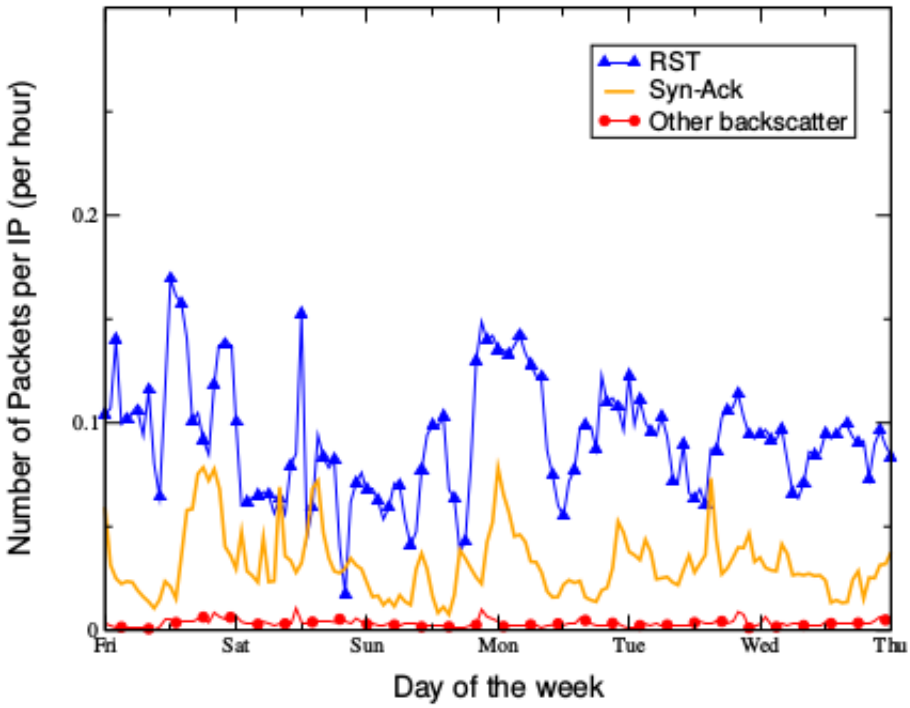


Figure 1: Packets per hour by type over one week in an unused /8 (Pang *et al.*, 2004)

Important advances in detecting and characterising IBR include determination of packet sources (as packet source addresses can be forged – this is called *spoofing*) (Barford *et al.*, 2006) and various advances in darknet placement and configuration and the analysis of data collected by darknets (Bailey *et al.*, 2005, 2006).

Beginning in October 2008 and continuing through 2009 and 2010, TCP packets destined for port 445 – products of the Conficker worm – became a large contributor to IBR, significantly increasing its prevalence (Wustrow *et al.*, 2010). This is still true today (Irwin, 2013).

Wustrow *et al.* (2010) discovered that modern IBR is increasingly made up of TCP SYN packets, with the total percentage of TCP SYN packets increasing from 62.7% of the total in 2006 to 93.9% in 2010, and decreasingly of the TCP SYN-ACK packets that formed the majority of the data recorded by Pang *et al.* (2004), which accounted for 26.1% of the total packets in 2006 but only 5.2% in 2010. Conficker's prevalence has also homogenised disparate IBR datasets to a much greater degree than discovered by Pang *et al.* (2004) (Wustrow *et al.*, 2010).

The homogenisation of IBR data in the time since Pang *et al.* (2004) is further corroborated by Irwin (2013), who discovered a significant degree of similarity in five /24 blocks occupying far apart areas in IPv4 address space, even extending to packets not characterised as products of Conficker. Nkhumeleni (2014) conducted further research into correlations between the datasets of these five /24 blocks, noting that ICMP packet traffic in the datasets showed similarities regardless of their network block locations. After removing the Conficker data (all packets destined to TCP port 445) Nkhumeleni (2014) found a strong cross-correlation between the African network blocks surveyed and a moderature cross-correlation between the North American network blocks surveyed.

Of some interest for the future is the advent of IPv6 IBR. At the moment it accounts for a very small percentage of both total IBR and total IPv6 traffic and has been characterised as largely caused by equipment misconfiguration, but this may change as IPv6 sees more wide-scale adoption (Czyz *et al.*, 2013).

## 2.1 Components of Internet Background Radiation

The four main components of IBR can be broadly categorised into two groups: active and passive traffic. Active traffic encompasses worms and scanning activities, both of which seek a response from the address they are destined for and are often malicious. Passive traffic encompasses DDoS backscatter and misconfigurations, which are merely the end results of DDoS attacks and network device misconfiguration respectively and carry no expectation of responses from other devices on the network (Irwin, 2011).

### 2.1.1 Worms

A worm is a form of virus that is network-aware and self-propagates across a network, using infected devices as springboards from which to infect other devices, often at random (Staniford *et al.*, 2002; Zou *et al.*, 2005). As a result of this random and rapid propagation, worms contribute significantly to IBR by scanning random and potentially

unused addresses in search of vulnerable systems. The worm scanning activity that contributes to IBR is largely on the TCP protocol, but a notable exception to this is the SQL Slammer worm, which scans the UDP 1434 port (Wustrow *et al.*, 2010).

Staniford *et al.* (2002) profiles the Code Red I, Code Red II and Nimda worms. Code Red I and II both self-propagated by spreading to randomly generated IP addresses, although notably a bug in Code Red I's random generation meant that the same seed was used for all random generation and thus the worm was not able to spread as far as it had presumably been designed to.

Code Red II, a worm that exploited the same Microsoft IIS web server vulnerability (CVE-2001-0500[1]) as Code Red I but was otherwise unrelated, did not share this defect and was able to propagate far more successfully. Code Red II selected IPs to attempt to spread to in the following manner: four out of eight times, it would choose an address from its host's /8 address space, three out of eight times, it would choose an address from its host's /16 address space, and the remaining one out of eight times, it would choose an address from the entire IPv4 address space. This allowed the worm to quickly spread across internal network, taking advantage of the likelihood that hosts with close IPs tend to be close together within the network.

Staniford *et al.* (2002) noted the difficulty of analysing the behaviour of Code Red II given that it was active at the same time as Code Red I, and as both made use of the same vulnerability and had generally similar behaviour, it was difficult to assign specific traces of Code Red behaviour to one or the other, especially if packet payloads were not taken into consideration. This is because the header data of packets from Code Red I and Code Red II would both have the same source and destination ports and the same protocols, with the only obvious difference being that Code Red II packets target a wider range of addresses than Code Red I packets.

The first internet worm monitoring system was introduced by Zou *et al.* (2005). The system detected worm activity by identifying trends in network traffic and relating them to models of worm propagation.

The advent of the Conficker worm is primarily responsible for the significant increase in the percentage of IBR in internet traffic (Wustrow *et al.*, 2010). Conficker exploited a vulnerability in the Windows Server service which allowed for arbitrary code execution on receipt of a specially formed RPC message. The worm has infected between 7 and 15 million hosts and is still spreading (Shin & Gu, 2010).

Irwin (2012) conducted an analysis of the Conficker worm's evolution based on IBR packets collected during the worm's 2008 outbreak. The packets identified as resulting from Conficker's activities were mostly TCP packets destined for port 445 (used for Microsoft Active Directory and Windows shares)[2] sent from hosts running a Microsoft Windows family operating system and targetting certain ranges, which matched the nature of the worm's operations.

[1] `http://cve.mitre.org/cgi-bin/cvename.cgi?name=2001-0500`

[2] `https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.txt`

### 2.1.2 Scanning Activities

Port scanning is the process of discovering information about a network by probing its devices with packets. It is often carried out by attackers or worms. By discovering which ports are open, closed and filtered on various devices on a network, attackers can discover IP addresses and their associated MAC addresses on a network, as well as what services are running and what firewall rules are in place, giving them enough information to launch an attack optimised to take advantage of any existing vulnerabilities (Modi *et al.*, 2013).

A common goal for attackers is the creation of a botnet (Staniford *et al.*, 2002). A botnet is the collective term for all computers compromised and controllable by a single controllable entity over a network connection. Botnets are closely related to both worms and DDoS attacks, in that a common means of generating and extending a botnet is writing a worm that infects hosts with bot software, and a common use for botnets is performing large-scale DDoS attacks (Cooke *et al.*, 2005). Botnets also give their controllers access to sensitive data on victimised systems and allow for impersonation of users (Staniford *et al.*, 2002).

Wustrow *et al.* (2010) noted the increase of SSH scanning activity to the point of significantly contributing to IBR, starting in 2007. Concurrent with this was the rise of TCP port 23 (traditionally used for Telnet connections[3]) scanning, indicating increased attempts to discover backdoors installed by worms.

Work has been done on systems like BotHunter (Gu *et al.*, 2007) and BotFinder (Tegeler *et al.*, 2012) to discover and identify the malware family (a system of categorisation of malware types by shared behavioural patterns (Rieck *et al.*, 2008)) of bots by the network traffic they produce. BotFinder was able to identify botnet-infected systems and botnet activity without analysing the content of any of the packets in the inspected network traffic (Tegeler *et al.*, 2012).

Bou-Harb *et al.* (2013) proposed a system of detecting network scans with a focus on the targets of the scans rather than their sources. The claim made was that as it is sometimes not possible to determine the sources of scans, techniques relying on this determination were prone to being less effective. The Bou-Harb *et al.* (2013) system instead chooses to cluster similar scans together under the assumption they come from the same source as an alternative to finding their sources.

### 2.1.3 DDoS attack backscatter

In a DDoS attack, one or more clients send a large amount of non-productive packets to a host, with the intent of either crashing or slowing the software running on that host or overwhelming the host's physical resources such as processing power and memory. The packets sent in this kind of attack are often spoofed: that is, their source addresses have been fabricated and do not match the addresses of the systems they were sent from (Moore *et al.*, 2006).

[3] `http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml`

A TCP SYN packet created in this way will prompt a response to its source address in the form of a TCP SYN-ACK packet, as part of the three-way handshake (Cerf *et al.*, 1974). When a victimised host sends responses such as these to packets with spoofed IP addresses, those responses will go to the spoofed IP, which may be an address residing in unused address space. This is especially probable if the spoofing is done randomly rather than with the intention of implicating a third party in the attack (Moore *et al.*, 2006). These response packets are termed "DDoS backscatter". Figure 2 demonstrates this.

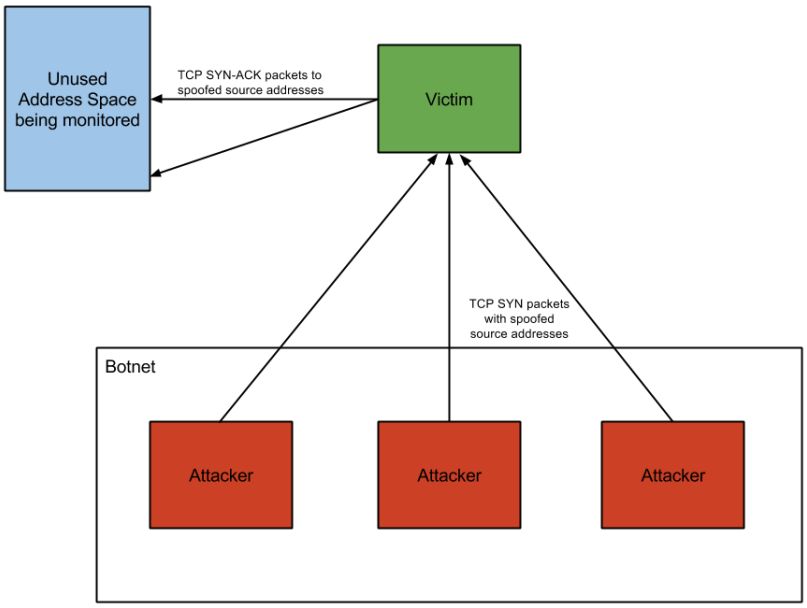


Figure 2: Diagram demonstrating DDoS backscatter

DDoS backscatter makes up a segment of IBR recorded by darknets and can be used to analyse DDoS activity. This technique was first used by Moore *et al.* (2006). The use of backscatter analysis in studying DDoS activity has been criticised by Mao *et al.* (2006), who cites a lack of address spoofing in most DDoS attacks conducted.

Fachkha *et al.* (2014) discusses a way of detecting DNS amplification–based DDoS activity from darknet data without relying on backscatter packets. This is discussed further in Section 4.2.2. The study highlights potential active components of IBR that are related to DDoS attacks but do not form part of passive backscatter.

### 2.1.4 Misconfigurations

Occasionally a network device on a host or a router will accidentally be set up to route packets to addresses within unused space monitored by a darknet. This creates small,

benign set of IBR data known as misconfigurations (Pang *et al.*, 2004). Misconfigurations account for the majority of IPv6 IBR at present (Czyz *et al.*, 2013).

The most common misconfigurations are unused addresses mistakenly entered in the address field, either in the systems themselves or in the Network Address Translation (NAT) routers (Irwin, 2011).

## 3 Border Gateway Protocol

A possible way of determining whether source addresses of packets in IBR are spoofed is to maintain a set of BGP routing tables for the data (Yao *et al.*, 2010). BGP routing is the system that governs how autonomous systems (ASes) route traffic to each other. This is shown in Figure 3. An AS, also called a routing domain, is a network with a clearly defined routing protocol between hosts. A single AS is often owned by an ISP.[4]

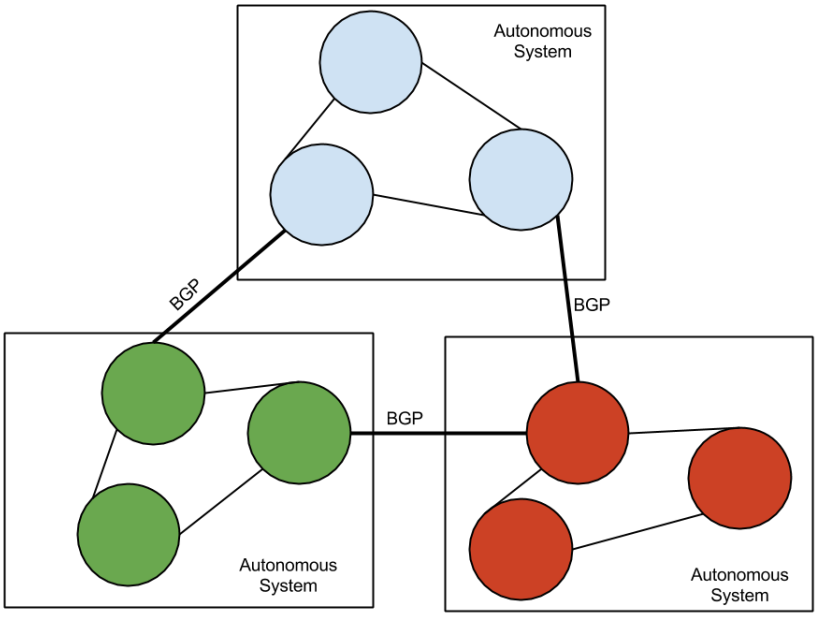


Figure 3: BGP routing between Autonomous Systems

BGP is an application layer protocol overlaying TCP. BGP messages (sent over TCP port 179) generally specify a network, a subnet and a number of attributes, the most relevant of which is the AS-path attribute, which indicates the order in which a packet must be transferred from AS to AS in order to reach the specified subnet of the specified network. Routers receive BGP route advertisement messages, prepend their own AS numbers to them, and then send those messages out as route advertisements to other routers (Schluting, 2006).

[4] `http://searchnetworking.techtarget.com/definition/autonomous-system`

Generally, implementations of BGP routing include mechanisms for *filtering* (Schluting, 2006) and *route flap damping* (Villamizar *et al.*, 1998). *Filtering* allows routers to reject route advertisements from unexpected sources. *Route flap damping* causes routers to ignore a route that disappears and reappears suddenly two or more times in quick succession (called *flapping*), with the time period in which to ignore the offending prefix increasing exponentially every time the route flaps.

Traditionally, BGP routing has few internal security features. All announcements of routes by all ASes are regarded as true. This raises security concerns, as BGP routes are trivially spoofed in the current system. Research has been done to elucidate the weaknesses of BGP routing and to attempt to make it more secure.

One of the earliest attempts at securing BGP was the Secure Border Gateway Protocol (S-BGP) (Kent *et al.*, 2000), a set of cryptographic attestations by which ASes could broadcast their veracity. More analytically, Qiu *et al.* (2007) discovered that legitimate BGP routes had stable historical structures and thus spoofing could be identified by comparing the current state of a routing system with its historical state.

Song *et al.* (2013) showed that S-BGP and similar systems do not successfully secure BGP against spoofing, as attackers can spoof routes indirectly through fundamental vectors S-BGP was not designed to protect against. The highly configurable nature of BGP implementations means that each AS can have different policies for route flap damping, filtering and the Minimum Route Advertisement Interval (MRAI) and thus an entire route can be invalidated by containing just one AS with badly configured policies (either due to ignorance or malicious intent).

Although it is widely known that BGP routing is vulnerable to attacks and many secure variations have been proposed, none have seen wide-scale adoption (Chan *et al.*, 2006).

# 4 Network Traffic Collection, Characterisation and Classification

The study and collection of IBR is a richly researched area that can be broadly split into two parts: the creation and refinement of tools for collecting IBR data, and the analysis of and extraction of meaning from the IBR data collected.

## 4.1 Collection of IBR data

IBR traffic is recorded via network telescopes, which are systems that monitor areas of unused address space and record packets sent to them (Moore *et al.*, 2004). Before Pang *et al.* (2004), network telescopes were used to study DDoS backscatter (Moore *et al.*, 2006) and worm activity (Moore *et al.*, 2003), but no overarching characterisation was done on the data.

Network telescopes variants are called darknets, greynets and other names such as dimnets (Irwin, 2011) based on the amount of unused addresses within or "darkness" of the address space they are set up to monitor.

### 4.1.1 Darknets

A darknet is a system that monitors traffic towards unused address space within a network (Moore *et al.*, 2004). A darknet is hosted on a network address that does not send any packets or interact with the outside network. Darknets set up to monitor small, well-defined address spaces will send Address Resolution Protocol (ARP) replies to routing requests for unused addresses in the network space. When deploying darknets to monitor larger address spaces (thousands of addresses and above), a router will route the entire unused address block to the darknet instead (Bailey *et al.*, 2006).

Legitimate traffic will not be captured by a darknet, as it has no reason for accessing addresses not set up to receive any traffic, and so all darknet traffic can be presumed to be unproductive and part of IBR (Bailey *et al.*, 2006). Thus there is no need to identify and separate out legitimate traffic from illegitimate and possibly malevolent traffic when analysing IBR.

Care needs to be taken to ensure that darknets are correctly positioned and configured so that they can record datasets of optimal usefulness. For example, darknets that monitor an area close to live address space received significantly more packets than ones further away (Bailey *et al.*, 2006).

### 4.1.2 Greynets

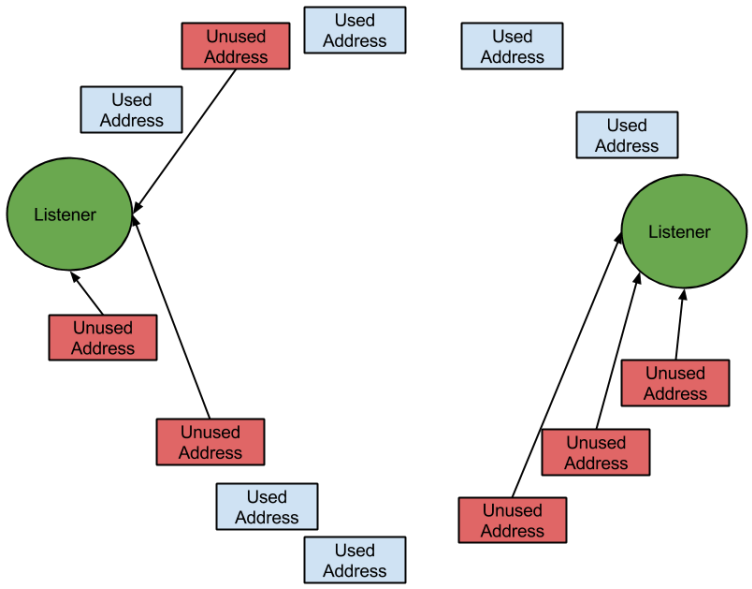


Figure 4: Diagram of a greynet with two listeners deployed on a small sample network

A greynet is a variant of a network telescope that, instead of monitoring a large block of contiguous unused address space like a darknet, monitors unused addresses scattered amongst a populated area of address space, within an organisational network for example. A greynet can thus be deployed on an enterprise network for the purpose of more direct monitoring than is possible with a darknet (Harrop & Armitage, 2005).

Thusly, a greynet can be deployed in a similar manner to the deployment of a darknet over a small address space as discussed above. A greynet can have many "listeners" – systems that receive unsolicited requests – spread out over the network. This is shown in Figure 4.

The only difference is that a greynet's listeners will send ARP replies to significantly less than one hundred percent of the routing requests to the address space.

## 4.2 Analysis of IBR data

The use of statistical analysis to characterise packet data in a network is a well-explored area. Paxson (1994) constructed analytical models of network flows created by applications using data such as packet length and flow duration.

Lin *et al.* (2009) expanded on the work done by Paxson (1994) by included application protocol-level data in the analysis conducted. Lin *et al.* (2009) observes that that packet size distribution (PSD) varies greatly between applications but is generally similar within applications. This provides a vector by which to cluster packet data according to application.

Barford *et al.* (2002) used signal analysis to identify anomalies in general network traffic collected on a border router at the University of Wisconsin and classify them into "long lived" (mass downloading of popular new software) and "short lived" (network outages, attacks and measurement anamolies) events.

### 4.2.1 Packet-level analysis

The most obvious way of analysing IBR datasets is to look at the trends in packet composition over time. Packet source address, destination address, protocol (usually TCP or UDP), source port and destination port are among the most commonly graphed data. Much can be determined about activity occurring on the network just by identifying trends in these data – how many unique source and destination addresses exist, the number of similar packets sent to each port and address, diurnal changes in packet behaviour and so forth.

For example, Conficker activity is evidenced by a large number of TCP SYN packets with destination port 445, occuring within the 1/4 of Internet address space Conficker propagation is limited to as a result of a bug in its psuedo-random number generation (Wustrow *et al.*, 2010). Another indicator of Conficker activity is that whereas most legitimate TCP connections are initiated by host sending three SYN packets (in case of data loss), Conficker only sends one or two (Aben, 2008).

DDoS activity can be inferred from large volumes of TCP SYN-ACK packets, with each single source being DDoS victim responding to TCP SYN packets with spoofed IP

addresses, as discussed in Section 2.1.3.

### 4.2.2 Network Flows

A network (or "message") flow is a set of packets sent between a given source address and a given destination address and port (Kerr & Bruins, 2001). Network flow identification is used by routers to determine correct port numbers for routing packets and to control access to them. Network flows can also be used for traffic analysis, providing a level of detail between general Simple Network Protocol (SNMP) statistics and highly detailed packet-level data (Sommer & Feldmann, 2002).

Network flows have been used to analyse traffic for cyber-defence purposes, i.e. identifying scanning, worm and DDoS attack activity (Chickowski, 2013; Yurcik, 2005). An example of this is Fachkha *et al.* (2014), which proposed a method of inferring DNS amplification–based DDoS attack activity by analysing DNS netflows discovered in darknet data.

DNS queries in the darknet space surveyed by Fachkha *et al.* (2014) were classifiable in three categories: DNS queries with spoofed source addresses sent by an attacker (the spoofed addresses being the victim's address), DNS queries sent by a compromised victim controlled by an attacker, and DNS queries sent as a scanning activity to infer the locations of open DNS resolvers. The sources of DDoS activity were to be inferred by the source addresses of the third type of DNS flow, wherein the source addresses of the packets would not be spoofed as the attacker would need to receive replies in order successfully carry out scanning activity.

### 4.2.3 Honeynets

A honeynet is a system of hosts purposefully made vulnerable and exposed to worms for the purpose of studying the worms.[5] Honeynets are often used in conjunction with darknets in systems for detecting malware and attacks. In this kind of system, a darknet or system of darknets will identify new threats in the traffic gathered and proxy the appropriate traffic to the system of honeynets for more in-depth analysis (Bailey *et al.*, 2005).

The principle difference between darknets and honeynets is that whereas darknets just record packet data, honeynets send appropriate responses to received packets in order to study their activities further. Darknets just record packet header data, whereas honeynets allow for study of actual packet payloads.

Yegneswaran *et al.* (2005) developed a system for providing network security personnel with in-depth network situational awareness (summarised, accurate data about moment-to-moment happenings within the network) using a honeynet system deployed on unused address space. The system highlighted two kinds of events: the advent of activity not previously seen on the network, and atypically large spikes of previously seen activities.

[5] `http://searchsecurity.techtarget.com/definition/honeynet`

## 5 Summary

The study and automated analysis of Internet Background Radiation can provide network administrators and network security experts with valuable intelligence on the potentially malicious scanning and DDoS activities occurring in and adjacent to their networks. It is also a useful dataset for the study of worm and botnet propagation.

As IBR is packet data sent to unused addresses, it includes no legitimate connections or service traffic and individual packets or netflows can all be assumed to result from worm activity, botnet and reconnaisance scanning activity, DDoS activity and occasionally the activity of misconfigured network devices.

The Border Gateway Protocol provides a method for routing packets between Autonomous Systems, and live BGP tables can be incorporated into IBR monitoring systems to identify spoofed IP addresses.

IBR is monitored using darknets and greynets. Activities in the wider internet can be inferred through analysing the packet data and network flows occurring within these datasets, and by sending data seen as "interesting" (novel or unusual in some manner) to honeynets for analysis of its packet payloads (Yegneswaran *et al.*, 2005).

As more sophisticated internet worms, botnets and DDoS attack techniques are developed, and as active worms such as Conficker continue to spread, the volume of IBR will continue to increase. By studying and characterising this data, researchers and security professionals have and can continue to develop robust warning systems and countermeasures to these threats.

## References


Aben, Emile. 2008. *Conficker/Conflicker/Downadup as seen from the UCSD Network Telescope.*

Bailey, Michael, Cooke, Evan, Jahanian, Farnam, Provos, Niels, Rosaen, Karl, & Watson, David. 2005. Data Reduction for the Scalable Automated Analysis of Distributed Darknet Traffic. *Pages 21–21 of: Proceedings of the 5th ACM SIGCOMM Conference on Internet Measurement.* IMC '05. Berkeley, CA, USA: USENIX Association.

Bailey, Michael, Cooke, Evan, Jahanian, Farnam, Myrick, Andrew, & Sinha, Sushant. 2006. Practical darknet measurement. *Pages 1496–1501 of: Information Sciences and Systems, 2006 40th Annual Conference on.* IEEE.

Barford, Paul, Kline, Jeffery, Plonka, David, & Ron, Amos. 2002. A signal analysis of network traffic anomalies. *Pages 71–82 of: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment.* ACM.

Barford, Paul, Nowak, Rob, Willett, Rebecca, & Yegneswaran, Vinod. 2006. Toward a model for source addresses of internet background radiation. *In: Proc. of the Passive and Active Measurement Conference.*

Bou-Harb, Elias, Debbabi, Mourad, & Assi, Chadi. 2013. A systematic approach for detecting and clustering distributed cyber scanning. *Computer Networks*, **57**(18), 3826–3839.

Cerf, V., Dalal, Y., & Sunshine, C. 1974 (Dec.). *Specification of Internet Transmission Control Program*. RFC 675.

Chan, Haowen, Dash, Debabrata, Perrig, Adrian, & Zhang, Hui. 2006. Modeling Adoptability of Secure BGP Protocol. *SIGCOMM Comput. Commun. Rev.*, **36**(4), 279–290.

Chickowski, Ericka. 2013. Using NetFlow Data For More Robust Network Security. Online. Available from: `http://www.darkreading.com/attacks-breaches/using-netflow-data-for-more-robust-network-security/d/d-id/1141085?` [Accessed on 16 May 2014].

Cooke, Evan, Jahanian, Farnam, & McPherson, Danny. 2005. The zombie roundup: Understanding, detecting, and disrupting botnets. *Page 44 of: Proceedings of the USENIX SRUTI Workshop*, vol. 39.

Czyz, Jakub, Lady, Kyle, Miller, Sam G, Bailey, Michael, Kallitsis, Michael, & Karir, Manish. 2013. Understanding IPv6 internet background radiation. *In: (to appear) Proceedings of the 13th ACM SIGCOMM Conference on Internet Measurement (IMC13), Barcelona, Spain*.

Fachkha, Claude, Bou-Harb, Elias, & Debbabi, Mourad. 2014. Fingerprinting Internet DNS Amplification DDoS Activities. *Pages 1–5 of: New Technologies, Mobility and Security (NTMS), 2014 6th International Conference on*.

Fuller, V., Li, T., Yu, J., & Varadhan, K. 1993 (Sept.). *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*. RFC 1519 (Proposed Standard). Obsoleted by RFC 4632.

Gu, Guofei, Porras, Phillip, Yegneswaran, Vinod, Fong, Martin, & Lee, Wenke. 2007. BotHunter: Detecting Malware Infection Through IDS-driven Dialog Correlation. *Pages 12:1–12:16 of: Proceedings of 16th USENIX Security Symposium on USENIX Security Symposium*. SS'07. Berkeley, CA, USA: USENIX Association.

Harrop, W., & Armitage, G. 2005 (Nov). Defining and Evaluating Greynets (Sparse Darknets). *Pages 344–350 of: Local Computer Networks, 2005. 30th Anniversary. The IEEE Conference on*.

Irwin, B. 2012 (Aug). A network telescope perspective of the Conficker outbreak. *Pages 1–8 of: Information Security for South Africa (ISSA), 2012*.

Irwin, B. 2013 (June). A baseline study of potentially malicious activity across five network telescopes. *Pages 1–17 of: Cyber Conflict (CyCon), 2013 5th International Conference on*.

Irwin, Barry Vivian William. 2011. *A framework for the application of network telescope sensors in a global IP network*. Ph.D. thesis, Rhodes University.

Kent, S., Lynn, C., & Seo, K. 2000. Secure Border Gateway Protocol (S-BGP). *Selected Areas in Communications, IEEE Journal on*, **18**(4), 582–592.

Kerr, D.R., & Bruins, B.L. 2001 (June 5). *Network flow switching and flow data export*. US Patent 6,243,667.

Lin, Ying-Dar, Lu, Chun-Nan, Lai, Yuan-Cheng, Peng, Wei-Hao, & Lin, Po-Ching. 2009. Application classification using packet size distribution and port association. *Journal of Network and Computer Applications*, **32**(5), 1023 – 1030. Next Generation Content Networks.

Mao, Z. Morley, Sekar, Vyas, Spatscheck, Oliver, van der Merwe, Jacobus, & Vasudevan, Rangarajan. 2006. Analyzing Large DDoS Attacks Using Multiple Data Sources. *Pages 161–168 of: Proceedings of the 2006 SIGCOMM Workshop on Large-scale Attack Defense*. LSAD '06. New York, NY, USA: ACM.

Modi, Chirag, Patel, Dhiren, Borisaniya, Bhavesh, Patel, Hiren, Patel, Avi, & Rajarajan, Muttukrishnan. 2013. A survey of intrusion detection techniques in Cloud. *Journal of Network and Computer Applications*, **36**(1), 42 – 57.

Moore, D., Paxson, V., Savage, S., Shannon, C., Staniford, S., & Weaver, N. 2003. Inside the Slammer Worm. *IEEE Security and Privacy*, **1**(4), 33–39.

Moore, David, Shannon, Colleen, Voelker, Geoffrey M, & Savage, Stefan. 2004. *Network telescopes: Technical report*. Department of Computer Science and Engineering, University of California, San Diego.

Moore, David, Shannon, Colleen, Brown, Douglas J., Voelker, Geoffrey M., & Savage, Stefan. 2006. Inferring Internet Denial-of-service Activity. *ACM Trans. Comput. Syst.*, **24**(2), 115–139.

Nkhumeleni, Thizwilondi Moses. 2014. *Correlation and Comparative Analysis of Traffic Across Five Network Telescope*. Masters thesis, Rhodes University.

Pang, Ruoming, Yegneswaran, Vinod, Barford, Paul, Paxson, Vern, & Peterson, Larry. 2004. Characteristics of Internet Background Radiation. *Pages 27–40 of: Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement*. IMC '04. New York, NY, USA: ACM.

Paxson, Vern. 1994. Empirically Derived Analytic Models of Wide-area TCP Connections. *IEEE/ACM Trans. Netw.*, **2**(4), 316–336.

Qiu, Jian, Gao, Lixin, Ranjan, Supranamaya, & Nucci, Antonio. 2007. Detecting bogus BGP route information: Going beyond prefix hijacking. *Pages 381–390 of: Security and Privacy in Communications Networks and the Workshops, 2007. SecureComm 2007. Third International Conference on*. IEEE.

Rieck, Konrad, Holz, Thorsten, Willems, Carsten, DÃŒssel, Patrick, & Laskov, Pavel. 2008. Learning and Classification of Malware Behavior. *Pages 108–125 of:* Zamboni, Diego (ed), *Detection of Intrusions and Malware, and Vulnerability Assessment.* Lecture Notes in Computer Science, vol. 5137. Springer Berlin Heidelberg.

Schluting, Charlie. 2006. Networking 101: Understanding BGP Routing. Online. Available from: `http://www.enterprisenetworkingplanet.com/netsp/article.php/3615896/Networking-101-Understanding-BGP-Routing.htm`. [Accessed on 22 May 2014].

Shin, Seungwon, & Gu, Guofei. 2010. Conficker and Beyond: A Large-scale Empirical Study. *Pages 151–160 of: Proceedings of the 26th Annual Computer Security Applications Conference.* ACSAC '10. New York, NY, USA: ACM.

Sommer, Robin, & Feldmann, Anja. 2002. NetFlow: Information Loss or Win? *Pages 173–174 of: Proceedings of the 2Nd ACM SIGCOMM Workshop on Internet Measurment.* IMW '02. New York, NY, USA: ACM.

Song, Yang, Venkataramani, A., & Gao, Lixin. 2013 (July). Identifying and Addressing Protocol Manipulation Attacks in 'Secure' BGP. *Pages 550–559 of: Distributed Computing Systems (ICDCS), 2013 IEEE 33rd International Conference on.*

Staniford, Stuart, Paxson, Vern, Weaver, Nicholas, *et al.* 2002. How to Own the Internet in Your Spare Time. *Pages 149–167 of: USENIX Security Symposium.*

Tegeler, Florian, Fu, Xiaoming, Vigna, Giovanni, & Kruegel, Christopher. 2012. Botfinder: Finding bots in network traffic without deep packet inspection. *Pages 349–360 of: Proceedings of the 8th international conference on Emerging networking experiments and technologies.* ACM.

Villamizar, C., Chandra, R., & Govindan, R. 1998 (Nov.). *BGP Route Flap Damping.* RFC 2439 (Proposed Standard).

Wustrow, Eric, Karir, Manish, Bailey, Michael, Jahanian, Farnam, & Huston, Geoff. 2010. Internet Background Radiation Revisited. *Pages 62–74 of: Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement.* IMC '10. New York, NY, USA: ACM.

Yao, Guang, Bi, Jun, & Zhou, Zijian. 2010. Passive IP traceback: capturing the origin of anonymous traffic through network telescopes. *SIGCOMM Comput. Commun. Rev.*, **41**(4), –.

Yegneswaran, Vinod, Barford, Paul, & Paxson, Vern. 2005. Using honeynets for internet situational awareness. *In: In Proceedings of the Fourth Workshop on Hot Topics in Networks (HotNets IV).*

Yurcik, William. 2005. Visualizing NetFlows for Security at Line Speed: The SIFT Tool Suite. *Pages 169–176 of: LISA.*

Zou, C.C., Gong, W., Towsley, D., & Gao, Lixin. 2005. The monitoring and early detection of Internet worms. *Networking, IEEE/ACM Transactions on*, **13**(5), 961–974.