UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
INSTITUTO DE INFORMÁTICA
CURSO DE CIÊNCIA DA COMPUTAÇÃO

ARTHUR VINÍCIUS CUNHA CAMARGO

Exploring Network Telescopes in the Age of IPv4 Exhaustion

Work presented in partial fulfillment
of the requirements for the degree of
Bachelor in Computer Science

Advisor: Prof. Dr. Lisandro Granville
Coadvisor: Prof. Dr. Leandro Bertholdo

Porto Alegre
February 2024

ACKNOWLEDGEMENTS

I extend my deepest gratitude to my family for their unwavering support throughout my academic journey. Specifically, I want to acknowledge my mother, Eliane, my father, Márcio, and my sister, Cecília, for their profound impact on my existence and overall happiness. Furthermore, I am immensely thankful to my advisors, Leandro Bertholdo and Lisandro Granville, whose guidance, opportunities, and invaluable insights have been fundamental in the completion of this work. Their mentorship has not only shaped my academic growth but also changed my mindset towards science in general. Lastly, I want to express my appreciation to you, the reader. Your engagement with my work validates its purpose and significance. Without your interest and attention, the effort invested in its creation would be incomplete. Thank you for being an essential part of this journey.

ABSTRACT

Cyber threat intelligence relies on network telescopes for detecting attack, and emerging threats, traditionally utilizing a substantial portion of the IPv4 address space. However, the escalating scarcity and value of this resource force universities and companies to grapple with the challenge of re-purposing their address spaces, potentially impacting cybersecurity effectiveness and hindering research efforts. In this thesis we investigate the historical usage of IPv4 addressing space in network telescopes and explores the impact of reducing this space on their ability to identify attackers and collect valuable research data. We do explore two network telescopes with the intention to assess the number of unique sources a reduced version is able to capture and to find out if there are IPs that are preferred over others. Our findings reveal that even halving the allocated space for a network telescope may still permits the detection of 80% of unique cyber attack sources, and the address allocation schema have little to none influence in this detection.

Keywords: Darknet. Network-Telescope. Internet Measurement. Security.

Explorando Network telescopes na era da exaustão do IPV4

RESUMO

A inteligência de ameaças cibernéticas depende de "network telescopes"para detectar ataques e ameaças emergentes, tradicionalmente utilizando uma parte substancial do espaço de endereçamento IPv4. No entanto, a crescente escassez e valor desse recurso obrigam universidades e empresas a lidar com o desafio de redirecionar seu espaço de endereçamento, potencialmente impactando a eficácia em sua cibersegurança e prejudicando pesquisas. Nesse trabalho de conclusão, investigamos o uso histórico do espaço de endereçamento IPv4 em "network telescopes"e exploramos o impacto da redução desse espaço em sua capacidade de identificar atacantes e coletar dados. Nós exploramos dois network telescopes com a intenção de verificar o número de origens únicas que uma versão reduzia consegue capturar e descobrir se há IPs que são preferíveis do que outros. Nossas descobertas revelam que mesmo reduzindo pela metade o espaço alocado para um "network telescope", ainda é possível detectar 80% das fontes únicas de ataques cibernéticos, e o esquema de alocação de endereços tem baixa influência nessa detecção.

Palavras-chave: Darknets. Cyber-segurança. Medição da Internet. Network-Telescope.

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS AND ACRONYMS

IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
IANA	Internet Assigned Numbers Authority
RIR	Regional Internet Registries
LACNIC	Latin America and Caribbean Network Information Centre
APNIC	Asia-Pacific Network Information Centre
AFRINIC	African Network Information Centre
ARIN	American Registry for Internet Numbers
RIPE NCC	Réseaux IP Européens Network Coordination Centre
IDS	Intrusion Detection System
IoT	Internet of Things
DDoS	Distributed Denial of Service
DoS	Denial of Service
IBR	Internet Backscatter Radiation
IMS	Internet Motion Sensor
CAIDA	Center for Applied Internet Data Analysis
DHCP	Dynamic Host configuration Protocol

CONTENTS

1 INTRODUCTION

Network telescopes, also known as “darknets” (MOORE et al., 2004), capture and record unsolicited Internet traffic directed towards globally routed but unused IP address space. While network telescopes have been utilized for years, they remain essential tools for detecting and studying cyber threats and global events.

The main application of network telescopes includes monitoring and analyzing Internet traffic, helping cybersecurity experts identify new threats, attack patterns, and understand the behavior of potential attackers. They have been used for years to observe cyberattacks on an Internet scale, such as botnets (ANTONAKAKIS et al., 2017), distributed denial of service (DDoS) (MOORE et al., 2004; JONKER et al., 2017), and network scan campaigns (RICHTER; BERGER, 2019a; CABANA et al., 2021), providing a myriad of insights on malicious, unwanted, and unexpected behavior of cyberattacks.

Network telescopes typically collect large volumes of data. For example, a /16 sensor can generate between 100GB and 1TB of data daily. Analyzing this massive amount of data is challenging, but the use of artificial intelligence (AI) techniques is transforming this scenario (D’ANDRÉA et al., 2023). Advanced IA methods significantly improve one’s ability to understand and interpret the data collected by those sensors, thereby providing deeper insights into emerging cyber threats.

Despite their many benefits, network telescopes also come with an inherent challenge: IPv4 address space is now a scarce and expensive resource. For example, an address space /19 (8,192 IP addresses) is rated between US$ 357,990 to US$ 395,673 (IPv4 Global, 2023). In this context, companies, research networks, and universities face a growing pressure to release their IPv4 address space used in those sensors in favor of other uses, or even to sell or rent these addresses.

Recent studies, such as the exploration of dynamic darknets in cloud environments (PAULEY; BARFORD; MCDANIEL, 2023), encounter financial challenges due to the rising costs of IPv4 address space. Beginning January 1, 2024, major cloud providers like AWS and Google Cloud will impose new charges for IPv4 usage (HUIDES; SANTHANAM; LEHWESS, 2023) (Google Cloud, 2023). These costs could present obstacles for ongoing and future telescope research in cloud platforms.

To address this situation, we conduct a two-part investigation focused on network telescopes. In the first part, we explore the state and size of network telescopes. In the second part, we examine the impact of reducing the size of an already small telescope

and how it affects its effectiveness in detecting cyber threats. To address these issues, we formulate two research questions that we intend to answer in this thesis:

1. **What is the typical range of address spaces utilized in network telescopes ?**
2. **What is the impact of reducing address space in network telescopes on the quality of cyber threat detection?**

To address the first research question, we conduct a meticulous literature review to determine if there has been a decline in the size of network telescopes and to assess its magnitude. Additionally, we provide insights into the primary objectives of each network telescope and highlight some of their unique features, including size.

Afterward, we initiate an exploration of two network telescope datasets to examine the impact of reducing their address space, addressing the second research question. Our analysis involves exploring the effects of reduced IPv4 address space on traffic volume and unique source detection. Our objectives are to understand how this reduction will affect their effectiveness in threat detection and to identify an optimal approach for allocating the limited IPv4 address resources. This exploration considers various sampling techniques suggested for network telescopes, and contributes to a better understanding of network telescope dynamics, offering insights for addressing challenges posed by IPv4 scarcity in the future of this cyber threat intelligence technique.

This thesis is organized as follows: In chapter 2, we provide definitions and a background about network telescopes, how they are normally deployed and some common applications. Furthermore, chapter 3 shows a literature review addressing key concepts and new approaches in the area. In chapter 4, we detail the methods we use to review network telescopes and how we assess the impact of reducing its IPv4 address space. In chapter 5 we presents our findings, and in chapter 6 we review our research questions and summarize our main findings. Finally, in chapter 7 we provide some insights about what can be done in the future.

2 BACKGROUND

In this section we provide background theory of network telescopes and some of their main applications. The objective is to show the main concepts with the intention to provied the reader information about the topic.

2.1 Network Telescopes

A network telescope or “darknet” is a way of logging unused IP address space of the Internet. Due to the unadvertised nature of those “dark” spaces, all the traffic received in this infrastructure is unsolicited and very likely malicious with a very low number of false positives. The deployment of these systems are considered straightforward as they don’t normally respond to its initial requester. The term “network telescope” is known by various alternative terms such as darkspace, darknet and black hole monitors. Throughout this thesis, we will refer only to the term “network telescope”, or simply “telescope” to mantain terminology harmony.

Network telescopes are normally considered to be a passive method of monitoring, using a trap based approach which means it tricks and traps adversaries in order to collect malicious activity (FACHKHA; DEBBABI, 2016). Contrary to honeypots, where the source tries to simulate a live host with functionalities, the telescopes does not ever respond to the resquester and just logs them all. The main advantage of utilizing this technique over honeypots is its easier deployment and management. The latter can still be compromised if not carefully designed and requires more computational power. However, a drawback of network telescopes is their lack of interactivity with the target. Since they do not respond to requests, there is no possibility to explore further what a malicious actor could do next.

2.1.1 Deployment

The first step of darknet monitoring according to (FACHKHA; DEBBABI, 2016) is to deployment of the sensors. This step requires the configuration of DHCP or the upstream router to forward the packet to the sensor. A basic network telescope deployment can be seen in Figure 2.1. Depending on the size of the network telescope it can

generates between 1GB to 1TB of data daily. That means that processing and storing the information have to be considered, recent advancements in data analysis, automation techniques, and artificial intelligence have leveraged the usability of large datasets generated by network telescopes. This enhancement aims to boost their effectiveness in cyber threat analysis, resulting in increased interest in network telescopes in recent years.

Figure 2.1: Example of simple network telescope deployment

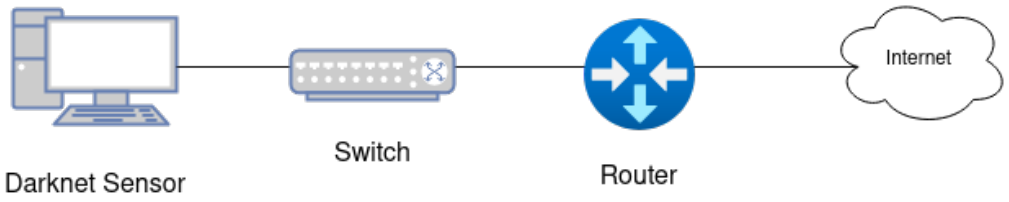

Source: Adapted from (FACHKHA; DEBBABI, 2016)

2.1.2 Applications

Network telescope sensors have various applications, including the analysis of Internet scan campaigns (RICHTER; BERGER, 2019a), locating botnets (MALÉCOT; INOUE, 2014), observing the proliferation of Internet worms (HARDER et al., 2006), and the analysis of Internet Backscattering Radiation (IBR) (BALKANLI; ZINCIR-HEYWOOD, 2014). These sensors are valuable for detecting and studying such threats, how they spread across the Internet and how attackers select their targets.

2.1.2.1 Identifying Internet Scan Campaigns

Internet scan campains are systematic efforts to scan large portions of the Internet for specific vulnerabilities, services, or devices. These campaigns are usually conducted by security researchers, cybersecurity professionals, or even malicious actors such as hackers. Those can normally be seen throught network telescopes as a scan not always discriminate between active or inactive hosts. It is possible to register scans even with a small number of addresses depending on the magnitude of the campaign, as the objective is to discover and search hosts systematically, scan campaigns normally tend to hit multiple addressess with the same type of resquest.

Not all scans are malicious, however. Some initiatives use scans to measure the Internet, such as Shodan and Cen++sys. They are generally considered benign scans because they are designed for legitimate purposes such as network monitoring, security

research, and reconnaissance. These platforms provide valuable information about devices connected to the internet and help organizations identify security vulnerabilities to improve their network security posture. Additionally, (SHODAN...,) and (CENSYS...,), allow individuals and organizations to request that their networks be excluded from scanning.

Network telescopes use to be effective on detecting scan campaigns. For instance, (RICHTER; BERGER, 2019a) utilizes a the Akamai/MIT network telescope that is composed of more than 89,000 CDN servers. The authors fixed in 100 the number of packets being sent by the same source in order to be considered a scan. Additionally, some timeouts was estipulated in order to gather scans with a temporal consistency. The paper shows that in total, they identified 2.2M scans, and which contain 87% of all logged traffic in their dataset.

2.1.2.2 Botnets Detection Using Network Telescopes

A botnet is a network of computers or Internet-connected devices that have been infected with malicious software, often without the owners' knowledge. These compromised devices, also known as "bots" or "zombies," are controlled remotely by some offender. Botnets are normally utilized to conduct malicious activities as: distributed denial of service attacks (DDoS), spam, phising and vulnerabilities scans campaigns. Because of their nature, some activities performed by those bots can be seem and tracked by network telescopes.

(MALÉCOT; INOUE, 2014) Tries to confirm data from a mysterious dataset about a year long scan of the Internet that was released anonymously in 2013, being called "Internet Census of 2012". The source of the data, claimed that it utilized 420 thousand compromised nodes (Carna Botnet) in order to make that scan. The paper utilizes a network telescope in order to infer the existence of that botnet, that used telnet brute-force attacks, and was able to recover some of the hosts that were compromised.

2.1.2.3 Observing Worm Proliferation

A worm is a type of malicious software (malware) that is capable of self-replication and spreading across networks without requiring human intervention. Unlike viruses, which typically require user interaction to spread, worms can propagate independently by exploiting vulnerabilities in network services or by using other means such as social en-

gineering. Because of its self-replication attribute, worms can also be seeing in network telescopes and sometimes even be estimated the time that new infections will occur and derive a sequence of infections.

(HARDER et al., 2006) Recreates some ofthe patterns and infection rules that were used by the Sasser worm in 2006 using a Class C network telescope. The paper is the first one to observer a long-range dependency in network traffic generated by malware. The authors concludes that it is possible to distinguish Internet noise and malware induced traffic.

2.1.2.4 Internet Backscattering Radiation

Internet Backscattering Radiation (IBR) is a side effect of using spoofed packets, such as those in DoS/DDoS attacks, where the attacker spoofs the IP address space from the network telescope, and the victim responds to that target. Backscatter traffic are one of the most common topic studied by large telescopes, as (CAIDA, 2024) and (Merit Network,). Because of their nature, it is harder to study them with smaller telescopes, as the number of reflected packets are very small comparated to the whole Internet noise that comes from scans, worm, spams and botnets.

(BALKANLI; ZINCIR-HEYWOOD, 2014) conducted a study on Internet Background Radiation (IBR) using CAIDA's network telescope over a period of 4 years. They also utilized 2 open source Intrusion Detection Systems (IDS) to understand if they could identify different attack behaviors in the traffic. Their study revealed that TCP-based attacks are by far the most common, and it remains very challenging for IDS and one-way traffic analysis tools to detect and understand backscattering traffic

3 RELATED WORKS

In this chapter, we delve into existing research concerning network telescopes. Our aim is to summarize findings and conclusions from other researchers, and to understand techniques used to enhance network telescopes for cybersecurity or to optimize IPv4 address space usage while maintaining its usefulness.

Despite the existence of large, well-known global network telescopes like CAIDA (CAIDA, 2024) and Merit (Merit Network,), many organizations face constraints in allocating addresses for threat intelligence. Deploying and maintain a network telescope is often hindered by the scarcity of unused IP address space, making it a significant entry-level barrier. Our research has noted a limited examination of the resources invested by companies and researchers in network telescopes. Thus, in section 4.1, we review existing literature to identify utilized network telescopes in research.

3.1 Network Telescope size

Several telescopes, including CAIDA and Merit, have reduced their size over the years. One of the earliest approach to employing a smaller address space for a network telescope while preserving the benefits for cyber threat intelligence was proposed by (HARROP; ARMITAGE, 2005). The authors suggested sparsely populating the sensors between actively used address spaces coining their approach as “greynets”— a mix of unused address space within specific subnets.

(HARROP; ARMITAGE, 2005) aimed to analyze attacks that tend to exploit topological neighborhoods, asserting that useful levels of network scan detection could be achieved with a smaller ’dark’ address space. However, they did not quantify the number of unique sources they might lose with this reduction in address space. We address this point in section 5.3.

Following the subject of comparing or reducing of address space usage of network telescopes, there are other authors who explores that topic.

(PEMBERTON; KOMISARCZUK; WELCH, 2007) explores sampling network telescopes and focus on the arrival density of backscatter radiation using a /16 network. The paper uses four schemes to slice its address space, horizontal, a contiguous /24 block, vertical, 1 address of each /24 and, Random-30 and Random-256, that respectively, chooses random addresses from 30 and 256 /24s respectively. Their work concludes that

deploying a number of random /32 networks across the telescope is the best way to predict backscatter radiation activity.

(CHINDIPHA; IRWIN; HERBERT, 2018) compares how different subnets behave on the matter of collecting IBR. The article computes the overlap of unique sources IPs across the whole network telescope. It concludes that lower /24 subnets do receive more unique origin hosts that the others, noting that 67% of the sources does only scan one IP of the sensor.

(SORO et al., 2019) compare 3 network telescopes with different sizes, /19, /15, and three /24 to assess how the size of each sensor influence the efficiency and detection of different types of events. The study presented evidence that the sources of traffic significantly vary based on the IP range and the size of the network telescope. The authors analyzed one week of data, aggregating it around autonomous systems (ASes). They demonstrated that reducing the network telescope by half minimally affects the visibility of network scans but results in different behavior in backscatter analysis when considering ASes. In section 5.3 we replicate part of their experiment on the same /19 telescope (Darknet-BR) to analyse the impact on reducing a network telescope IP address size.

3.2 New Artificial Intelligence Application on Telescope Datasets

The sheer volume of data collected by network telescopes presents a significant challenge for traditional analytical methods. Extracting meaningful patterns and identifying noteworthy events amidst the noise requires sophisticated algorithms and computational power. This is where Artificial Intelligence applications has made considerate progress, offering powerful tools for processing, classifying, and interpreting telescope datasets with unprecedented efficiency and accuracy.

(SHAIKH et al., 2018) used network telescopes and a model classifier to identify infected IoT devices in enterprises. The work makes use of the CAIDA dataset in order to train their model and classifies requests as scans, DoS and misconfigurations. They utilized Gradient Boosting and Random Forest classification algorithm that presented a high recall and precision.

(CABANA et al., 2021) utilized a combination of network telescope traffic analysis and artificial intelligence to analyze reconnaissance attack campaigns against industrial control systems, allowing an automatic determination of the threat level associated with each campaign.

(KALLITSIS et al., 2022) uses a unsupervised learning approach on the Merit Telescope to create clusters and detect different types of scan and attacks. Thay way, fine tunning the clusterring algorithms, they could find more than 200 groups of scanners with different characteristics and patterns. They found 70 Mirai-related cluster, identified 20 clusters associated with TCP/445 scan- ning, i.e., the SMB protocol. that normally is a great target for malware and ransomware, and furthermore, they identified a lot of “heavy scanners” that not always are nefarious (i.e. Shodan and Censys.io).

(SORO et al., 2020) introduced community detection algorithms applied to represent network telescope activity as a graph, grouping hosts infected by a botnet that is actively scanning the network in search of vulnerable services.

3.3 Recent Telescope solutions on IPv4 address scarcity

Recent developments addressing the scarcity of IPv4 address space for constructing new network telescopes propose other approaches while upholding their primary goals for cyber intelligence.

The “cloud-native Internet telescope” (PAULEY; BARFORD; MCDANIEL, 2023) suggests deploying short-lived telescopes on virtual machines within a cloud provider, leasing IPv4 address space, and releasing it after use. Their results indicate that optimal price performance per IPv4 is achieved in 8 minutes, and 90% of the steady-state traffic to a given IP address, compared with a regular network telescope, can be observed after only 72 minutes.

Another approach, called ”meta-telescope” (WAGNER et al., 2023), proposes identifying “unlikely to be used” address space in central points of the Internet (a.k.a Internet Exchanges) and capturing unsolicited traffic to this address space. In their research, they were able to capture unsolicited traffic for more than 350k /24 blocks in over 7k ASes.

4 METHODOLOGIES

In this section, we initially review the literature to identify all known network telescopes, along with their deployment characteristics, address space usage, and relevant research results based on each infrastructure. In the second part, we investigate the impact of reducing the IPv4 address space in an existing network telescope. This section delves into the methodology employed in both cases.

4.1 Network Telescopes over time

To understand the current landscape of IP address utilization in network telescopes, we conducted a literature review. Specifically, we gathered information on the types of sensors being used for research or production purposes, aiming to present a fresh perspective on their deployment trends over time. Additionally, we intend to show the status and information gathered by the projects and its approaches.

To achieve this objective, we selected data from research papers and significant projects related to network telescopes. Given our primary focus on studying address space usage, certain key characteristics are deemed essential for each subject under review. These include the number of IPv4 addresses utilized, the date of deployment, and the primary objectives pursued by the authors.

To make the data more accurate, we only select survey, reviews, essays, databases and papers related to network telescopes that are at least from year 2000. Non-relevant works were not selected (e.g., white paper, experimental studies, and reports lacking the information we are intending to collect) . To be included, the documents must prove to be informative and descriptive and meet one of the following criteria: (i) published by a respected organization with strong scientific endeavor or (ii) published in influential journals and conferences. In addition to those requirements, it is important that the source expresses the steps used to deploy and maintain the network telescope, as well as the results that were achieved during its activity.

To collect high-quality research and studies, we conducted a manual search for keywords related to each topic on Google Scholar. We used keywords such as "darknets," "network telescope," or "network blackhole" in our search. After reading the abstract of each article, we excluded those that were not eligible. Then, were selected relevant studies for further reading while discarding off-topic articles and papers.

4.2 Network Telescope reduction impact

To explore the implications of reducing the address space utilized in a Network Telescope, we conducted a comprehensive analysis with one month of data collected from two different sources at different times: the Japanese NICTER Darknet in 2018 (HAN et al., 2022) and the Darknet-BR from December 2023 to January 2024 (SORO et al., 2019).

So as to investigate if there are more voluminous addresses, we inspected and compared the number of requests each address received in both telescopes, which are explained in section 4.2.2. Our main focus was the capacity of the telescope to capture unique sources. That way, we calculated the expected value of unique sources for reduced versions of both NICTER and Darknet-BR in section 4.2.3. Finally, in section 4.2.4 , sampling strategies based on (PEMBERTON; KOMISARCZUK; WELCH, 2007) approach was implemented in order to understand how different allocations affect the sensors potential.

4.2.1 Network telescope Datasets

The Network Incident Analysis Center for Tactical Emergency Response (NICTER) darknet is an integration of large-scale network monitoring for the analysis of cyber threats, such as botnets or DDoS attacks. Its dataset encompasses information from eight sensors distributed worldwide, covering networks ranging from /20 to /17. The dataset archives one month of data from October 2018 and includes only TCP SYN packets.

The Darknet-BR is a Brazilian network telescope operating on a /19 IPv4 prefix over which we have full control, enabling us to perform more in-depth analysis on the captured packets. We used a dataset from December 14, 2023, to January 14, 2024, in our analysis. Table 4.1 shows the period during which each dataset was collected, the daily volume of collected data, and the address space size of each sensor.

4.2.2 Distribution per IP

For the purpose of further analysing the relation of address space and threat detection abilities, was made an examination of the number of requests received by each IP address in the network telescopes.

Table 4.1: Network Telescope Datasets from NICTER and Darknet-BR.

Sensor	Size	Volume
NICTER-A	/17	10GB per day
NICTER-B	/18	6GB per day
NICTER-C	/20	2GB per day
NICTER-D	/20	2GB per day
NICTER-E	/19	3GB per day
NICTER-F	/18	6GB per day
NICTER-G	/21	800MB per day
NICTER-H	/21	800MB per day
Darknet-BR	/19	3GB per day

Figure 4.1: Probability of k being seen in the reduced telescope S

$$P_k = |S|\frac{T_k}{|N|}$$

For validation purpose we apply our method in another network telescope. We select from a dataset provided by NICTER, a telescope with several IP spaces. After analysing all of its sensors and notice they presents similar behaviour, we select just the one for further comparison (sensor E). This sensor have the same size as Darknet-BR telescope. It worth to mention we just analyzed the number of scan events (TCP-SYN) and unique IP sources in this study. We do not consider UDP or ICMP data to keep the comparison possible–NICTER-E just provided TCP-SYN data.

4.2.3 Expected value of unique sources

In order to estimate the number of unique sources expected for allocating different sizes of network telescopes, we utilized a probabilistic approach based on the number of different destinations each source has. Considering N the set of addresses in the network telescope and S a subset of N in a way that $N \subseteq S$. Naming K as the set of source IPs, captured by the telescope and T_k the number of times that same origin address $k \in K$ appears. The probability of k being observed by the smaller version S considering an uniform distribution is given by Figure 4.1.

In addition, we grouped the number of IPs k that target the same number of T_k together as $G_t k$, as they provide roughly the same information for our model. That way, it is possible to deduct the expected number of unique sources that the reduced version of the sensor will capture utilizing the formula of expected value Figure 4.2.

Figure 4.2: Expected number of unique sources received by a sampled telescope

$$E[S] = \sum_{k}^{k \in K} min(P_k, 1) G_t k$$

As our works revolves around gathering unique sources, it is important to limit the probability function to not surpass 1, as it means counting the same attacker more than once. The only parameter related to the reduced network telescope in this case, is the size, as an uniform distribution is being considered, the formula is not dependent on the individual addresses being picked in the smaller version.

4.2.4 Sampling strategies

One way that organizations have to mitigate the problems of IPv4 exhaustion and scarcity is to reduce their telescopes. To do that, it faces the dilemma of how to divide the blocks. Said that, using the Darknet-BR, we considered two possible alternatives: (1) a reduction from a /19 (8,192 IPs) to a contiguous /20 (4,096 IPs), evaluating if there is differences between the first and second /20, or; (2) reducing it to several /24s by adopting a sample strategy proposed by (PEMBERTON; KOMISARCZUK; WELCH, 2007). While the first solution (1) is the simpler one, we consider to study the sampling solution (2) as a possibility to minimize the network telescope lost potential.

For the first study, we consider to analyze just the horizontal sampling from (PEMBERTON; KOMISARCZUK; WELCH, 2007), since vertical sampling we consider operationally unfeasible–when delegating a prefix we lose control over the that range. In horizontal sampling, we select half of the /24 blocks for the telescope, while the others will be assigned for users. As for the second study, we considered selecting alternating /24 blocks, in a way that they are equally spaced, trying to cover a larger area of addresses. Complementary, we observed individual blocks at specific locations within the network telescope, such as the beginning, end, and middle, exploring potential correlations with the findings of (HARROP; ARMITAGE, 2005), and (CHINDIPHA; IRWIN; HERBERT, 2018).

5 RESULTS

In this section, we present our results from research on identifying network telescope initiatives over the years and our investigation into the impact of reducing the address space of a currently operational network telescope.

5.1 Address space utilization on Network Telescopes over time

After conducting our paper review on network telescopes (see section 4.1), we summarize the main telescopes initiatives we identified in table 5.1. It is important to note that network telescopes in the cybersecurity industry typically do not publish their data, address space size, or even their existence to safeguard the secrecy of their initiatives. Consequently, we could not gather much information about those environments.

Additionally, most network telescope initiatives referred in research also does not disclose the address space they use. They justify this approach to avoid *adversary traffic*—when an attacker avoids scanning or using the network telescope address to avoid being identified. We have listed the address spaces that we could verify.

Table 5.1: Summary of Network Telescope projects referred from 2000 to present

IPv4 Addr.	Year	Name	Comments
50,331,648	2010	APNIC/ARIN	APNIC and ARIN collaborated on IBR research utilizing unallocated addresses 1/8, 50/8, and 107/8. This telescope had a lifespan of 1 week in 2006. (WUSTROW et al., 2010)
17,048,576	2001	Internet Motion Sensor	Arbor Networks and the University of Michigan project deploys sensors in diverse locations to enhance the diversity, sparsity, and size of a Network Telescope. The IMS initiative seems ending in 2004 and spanning into Merit Telescope. (COOKE et al., 2004b)
16,777,216	2005	MERIT	Merit Network Telescope used the 35/8 address from 2005 to 2018. After this date the Michigan University formalized the Orion telescope with a smallest address space. (Merit Network,)
16,777,216	2001	UCSD-CAIDA	The UCSD Network Telescope, a project from the University of San Diego/US was built on the globally routed 44/8 prefix (former AMPRNet) from 2001 to 2019. (CAIDA, 2024)

Continued on next page

Table 5.1 continued from previous page

IPv4 Addr.	Year	Name	Comments
12,582,912	2019	UCSD-CAIDA	The UCSD Network Telescope reduced it size from a /8 to a /9 and /10 network. (CAIDA, 2024)
∼2,000,000	2012	SWITCH	Collect data from the address space from multiple networks across Switzerland. (SWITCH,)
626,944	2004	Team Cymru	Multiple sensors deployed by the company Team Cymru. (CYMRU,)
524,288	2014	Farsight	Farsight's Network Telescope, now part of DomainTools, offers data through subscription. (Farsight Security,)
475,136	2018	ORION-MERIT	Michigan State University's project, known as the Observatory for Cyber-Risk Insights and Outages of Networks, focuses on Internet backscatter radiation. Designed and engineered with support from the US-NSF, it consists of 1,856 /24 subnets. (Merit Network,)
270,000	2005	NICTER	The Japanese organization NICTER (Network Incident Analysis Center for Tactical Emergency Response) integrates its network telescope for large-scale monitoring and analysis of cyber threats, including botnets and DDoS. (NICTER WEB,)
178,000	2018	MIT-Akamai	The first network telescope built over a Content Delivery Network (CDN) infrastructure. It is composed of two IPs on each of the 89,000 Akamai servers across the globe. (RICHTER; BERGER, 2019b)
131,072	2018	NL-Darknet	Network telescope maintained by SurfNET in the Netherlands. (SORO et al., 2019)
65,636	2019	HEAnet	Ireland's National Education and Research Network Telescope. (O'HARA, 2019)
65,536	2004	IUCC/IDC Telescope	The Israel InterUniversity Computation Center (IUCC) Network Telescope. (IUCC, 2024)
65,536	2006	Anonymous	The University of Wellington, NZ, utilizes an undisclosed /16 network telescope to test various address sampling strategies for measuring arrival density. (PEMBERTON; KOMISARCZUK; WELCH, 2007)
65,536	2021	Anonymous	An undisclosed enterprise network telescope identifies a specific stateless-scanning malware, and a response is forged to slow down the malware's propagation, deceiving botnet scanners. The research is being conducted in Germany. (GRIFFIOEN; DOERR, 2023)

Continued on next page

Table 5.1 continued from previous page

IPv4 Addr.	Year	Name	Comments
8,192	2018	BR-Darknet	A /19 network telescope in Brazil (used in this thesis). (SORO et al., 2019)
4,096	2017	JP-Darknet	Another /20 network telescope hosted in Japan. (LAGRAA; CHEN; FRANçOIS, 2019), (ZAKROUM et al., 2023) and (D'ANDRÉA et al., 2023)
4,096	2010	INRIA	French Telescope at INRIA's High Security Laboratory. (LAGRAA; CHEN; FRANçOIS, 2019), (ZAKROUM et al., 2023), (D'ANDRÉA et al., 2023) and (HOUMZ et al., 2021)
765	2017	IT-Darknet	Italian network telescope (SORO et al., 2019)
512	2009	Rhodes University	The first known network telescope in the AFRINIC Region. (IRWIN, 2011)
512	2013	KISTI	Science and Technology Security Center South Korea (KISTI), does provide 2 sensors with a /24 mask. (GADHIA et al., 2015)
256	2006	–	After 2006, numerous minor initiatives deployed temporary network telescopes with short lifespans (1-2 years), ranging from /28 to /24, for specific research. (NIRANJANA; KUMAR; SHEEN, 2019), (EZE; SPEAKMAN; ONWUBIKO, 2020), (FENG et al., 2013), (AHMED, 2010) and (AHMED; CLARK; MOHAY, 2008)

From the paper review we observed that the majority of significant network telescopes emerged between 2000 and 2007, a period characterized by fewer issues related to IPv4 allocation. The onset of IPv4 address exhaustion was first announced by the Regional Internet Registry (RIR) in Asia in 2011, followed by announcements from other RIRs in subsequent years (APNIC, ; RIPE, ; LACNIC, ; AFRINIC,).

The IPv4 exhaustion resulted in the absence of new relevant network telescope initiatives and even a reduction in existing network telescopes in recent years. For instance, UCSD-NT/CAIDA, which is part of the US Protected Repository for the Defense of Infrastructure against Cyber Threats (PREDICT) program and its successor, the Information Marketplace for Policy and Analysis of Cyber-risk & Trust (IMPACT) (University of California San Diego - Center for Applied Internet Data Analysis, 2018), saw a reduction in its IPv4 address space from a /8 to a /9 plus /10 in 2019.

In Figure 5.1, we can better visualize the decline in the utilization of public IPv4 addresses for network telescopes over the years. The figure is built from a more comprehensive list of publications we investigated, encompassing 28 network telescopes (blue

Figure 5.1: Tendency graph of IPv4 usage on Network Telescopes in each year, considering the first deployment of each one. Here, we omitted data from initiatives who reduced the size of their telescope (i.e., UCSD and MERIT).

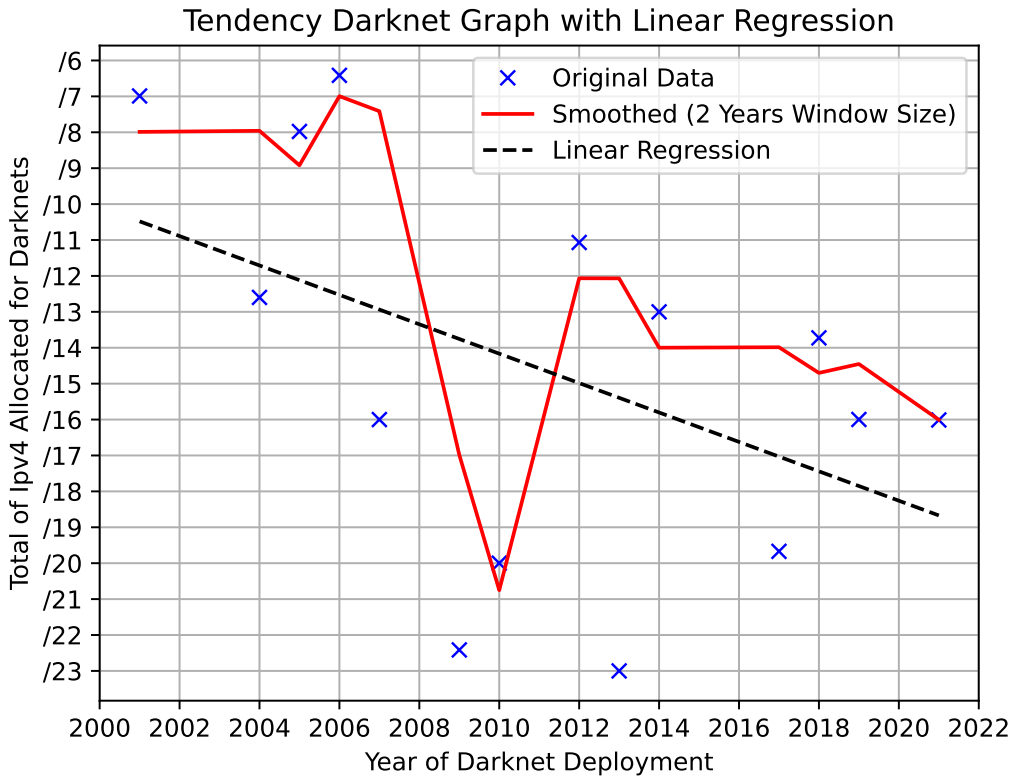

crosses). From this graph, is reasonable to infer an initial reduction around 2010, correlated with the depletion of IPv4 address in RIRs. The second reduction, around 2021 may be linked to the escalating prices of IPv4 in the market. Notably, during this period, portions of addresses from large network telescopes shifted to major companies such as Google and Amazon.

The reduction trend become more evident when we visualize the smoothed tendency over a two-year time window (red line). A linear regression (dashed back line) for this tendency also point to a possible deallocation of IPv4 address space for Network Telescopes after 2018– as observed in the case of UCSD-NT/CAIDA.

Given the difficulty of maintaining this address space active, new research initiatives such as *meta-telescope* (WAGNER et al., 2023) and *DScope* (PAULEY; BARFORD; MCDANIEL, 2023) aim to explore new ways to deploy temporary telescopes.

5.2 Analysing NICTER dataset

In this section, we observed some interesting characteristics in the IP distribution of each NICTER sensor. As pointed in subsection 4.2.2 we explored all the sensors of the Japanese telescope in order to find interesting caracteristics.

The data was presented as in Table 5.2, the sources was hashed to preserve the identity of the requesters, additionally the destination was also truncate in a way to just show the last 2 octect in order to preserve the anonymity of the sensors.

Table 5.2: Network Telescope NICTER Data types

Data type	Details
timestamp	received packet time (UNIX time)
hash[ip.src.upper16]	hash value of upper 16-bit source IP address
hash[ip.src.32]	hash value of 32-bit source IP address
ip.dst.lower16	lower 16-bit darknet destination IP address
tcp.dstport	16-bit TCP destination port number

Source: (NICTER WEB,)

The following figures show the number of requests that each destination address received in the 31 days of the sensor activity as commented on section 4.1. Exploring the sensors A, B and C we found that the allocations were not contiguous, contrary to the other ones, so we did a scalling that permitted to show all the resquests using a /16 mask in Figure 5.2, Figure 5.4 and Figure 5.6. All the other images shows only addresses that received requests. The x axis shows the IPv4 addresses in ascending order and the y axis the number of requests received by each destination.

It is possible to observe that there are some common addresses that normally receive around 100k requests and a second group of IPs that receive almost 200k and a third that receives around 250k. After some deeper analyzes we discovered that IPs ending in .1 and .2 receive the most requests, followed by the set from .3 to .25. That is a pattern that can be seen in all sensors and we mainly suspect telnet scans searching for IoT devices.

5.2.1 Sensor A

Sensor A, Figure 5.2 is by far the largest one we analyzed, it does contains some individual blocks that are not contiguous as it can be seen in Figure 5.2. This particular dataset, contains an outlier on the address 154.156 and that we believe that this is due to some mistakes and malfunctions that happened in the collection of the data. The re-

searchers that collected the dataset warned us about those types of problems in this sensor.

Figure 5.2: NICTER Sensor A, Requests per IPv4 (/17) inside a /16 frame

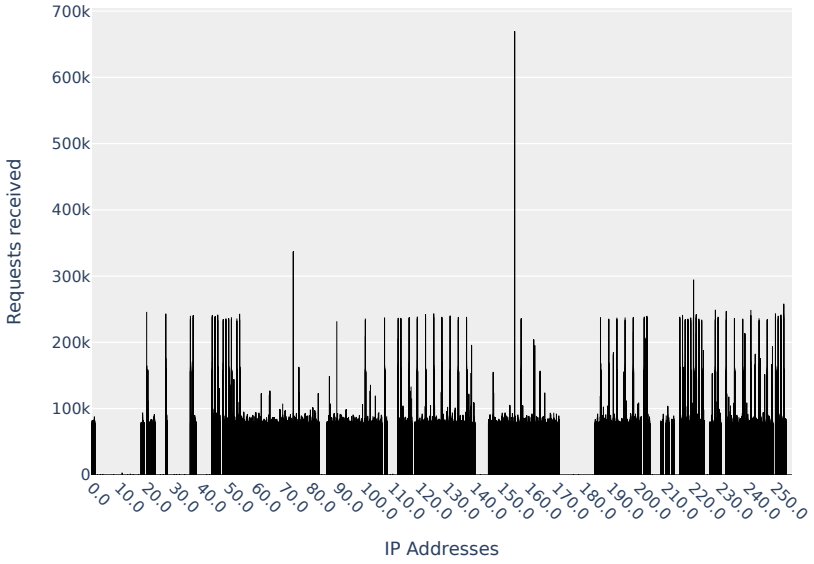

Figure 5.3: NICTER Sensor A, Requests per IPv4 (/17)

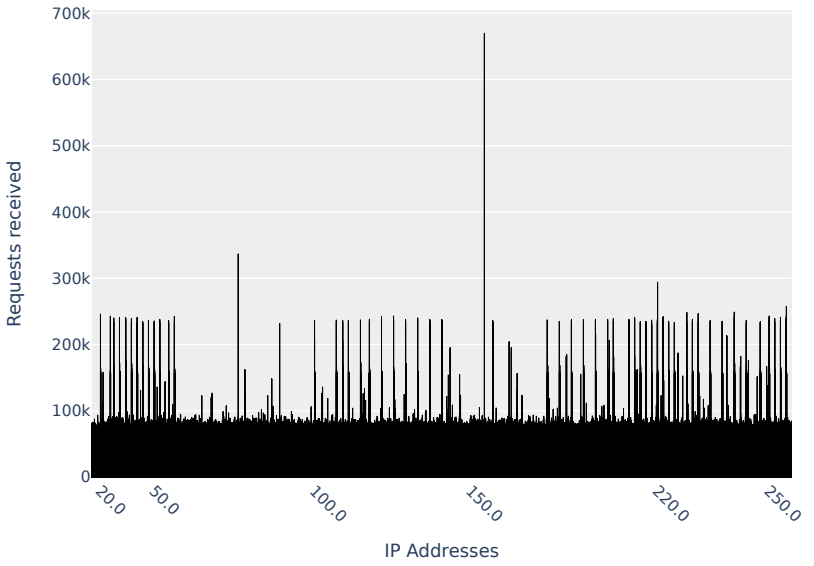

5.2.2 Sensor B

Sensor B, Figure 5.5 share some characteristics with the first one as it shows sparse IP blocks Figure 5.4. It is also possible to clearly see the request spikes also in the lower IP .1 and .2. This sensor also contains some outliers that received almost 400k requests.

Figure 5.4: NICTER Sensor B, Requests per IPv4 (/18) inside a /16 frame

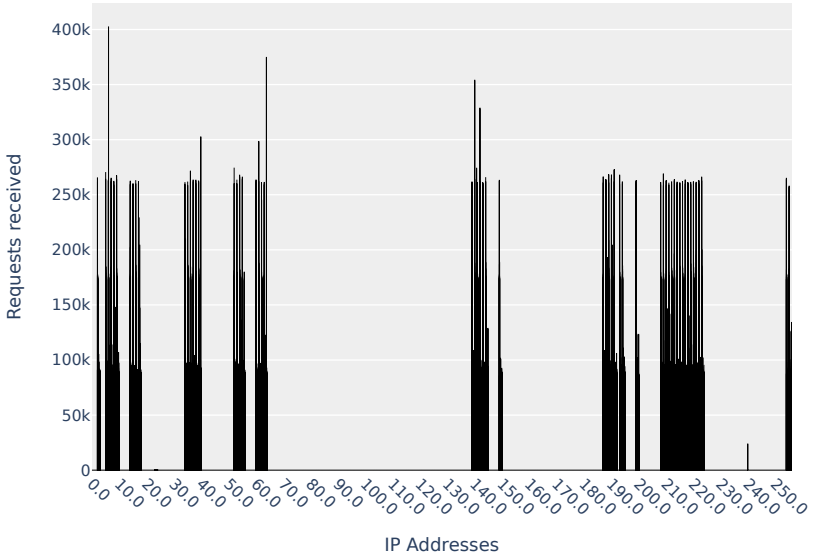

Figure 5.5: NICTER Sensor B, Requests per IPv4 (/18)

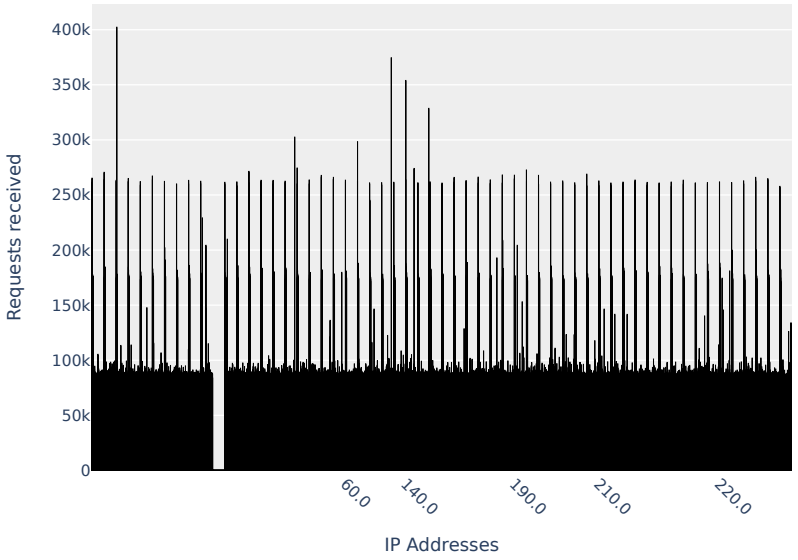

5.2.3 Sensor C

The sensor C, Figure 5.7 is much smaller than the previous ones, but still shows some non contiguous requests Figure 5.6. Here we can better see the second group of requests that are received by the x.3 - x.25 IPs. Those are also present in the bigger telescopes, but they do not appear as clearly because of the scale.

Figure 5.6: NICTER Sensor C, Requests per IPv4 (/20) inside a /16 frame

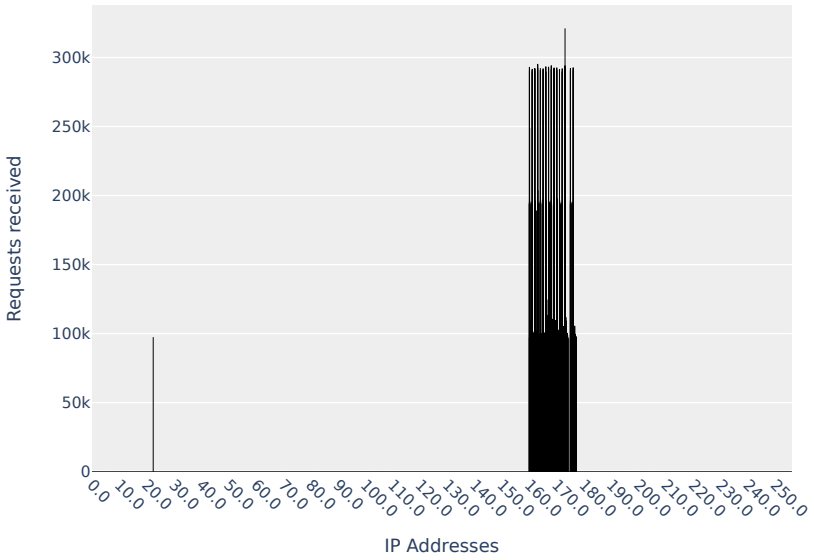

Figure 5.7: NICTER Sensor C, Requests per IPv4 (/20)

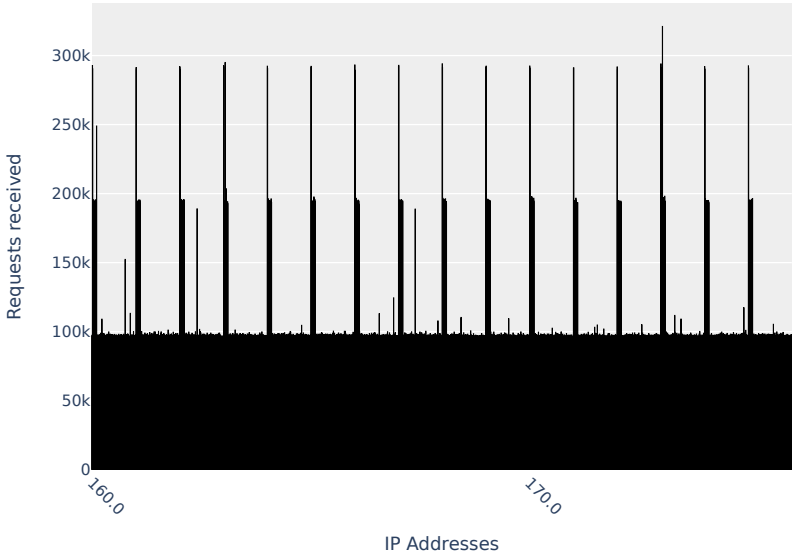

Figure 5.8: NICTER Sensor D, Requests per IPv4 (/20)

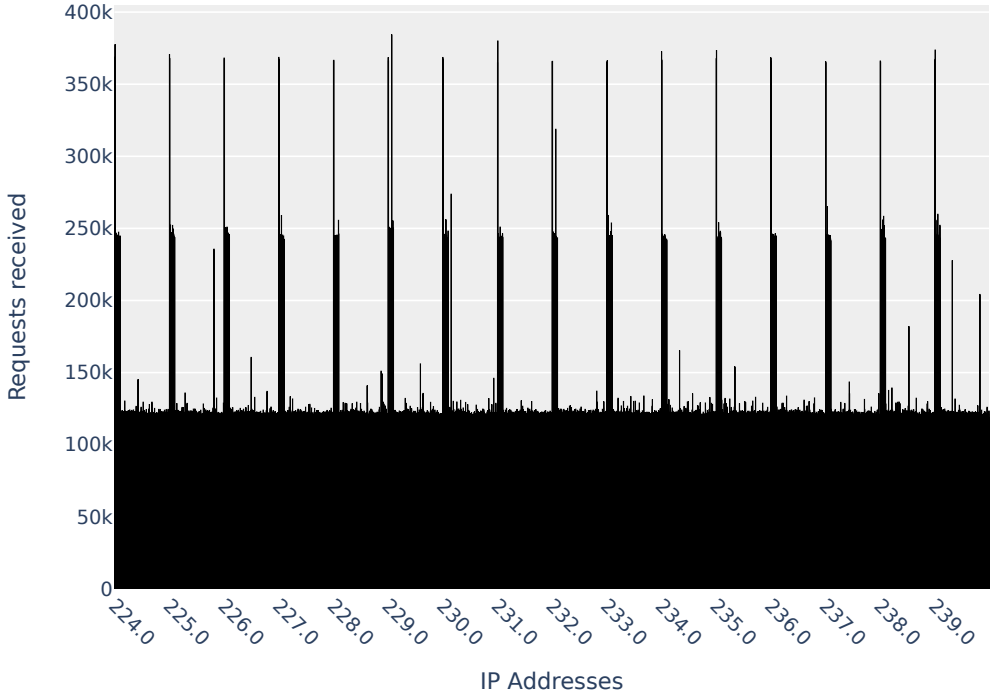

5.2.4 Sensor D

The sensor D, Figure 5.7 is the same size of sensor C, but all its addresses are contiguous. The same pattern appears again, a bigger number of requests appears in first part of each /24 block. Although, the number of overall requests received by this telescope was greater then the others. That is expected as the sensors are spread across the globe, so it is resonable to a sensor be more requested.

5.2.5 Sensor E

Sensor E, Figure 5.9 is the closest to our Darknet-BR in size (/19). That way, we compared their distribution on section 5.3. This sensor also reveals to be located in a lower part of the /16 block, as the numbering of IP destination suggests.

Figure 5.9: NICTER Sensor E, Requests per IPv4 (/19)

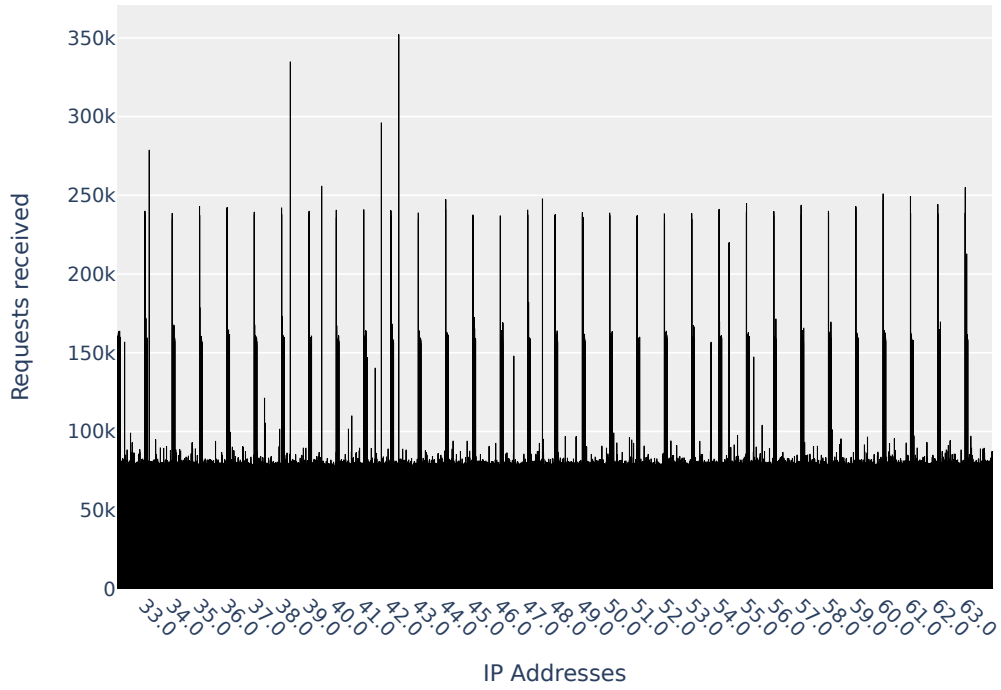

5.2.6 Sensor F

The sensor F, Figure 5.10 is also a continuous sensor and also shows the same patterns as the other ones. As the size of this sensor is the same as the size of subsection 5.2.2 it is possible to closely compare the distribution of both. Regardless of the continuity of sensor F in relation to subsection 5.2.2, the number of requests per IP looks very similar.

5.2.7 Sensor G

Sensor G, Figure 5.11 is among the smallest sensor, and because of that, it is possible to even more clearly see the request pattern that appears in the first 2 IP addresses, followed by the next 23. Similar to other sensors, as subsection 5.2.1 and subsection 5.2.2, there is an outlier that received more than 350k requests.

Figure 5.10: NICTER Sensor F, Requests per IPv4 (/18)

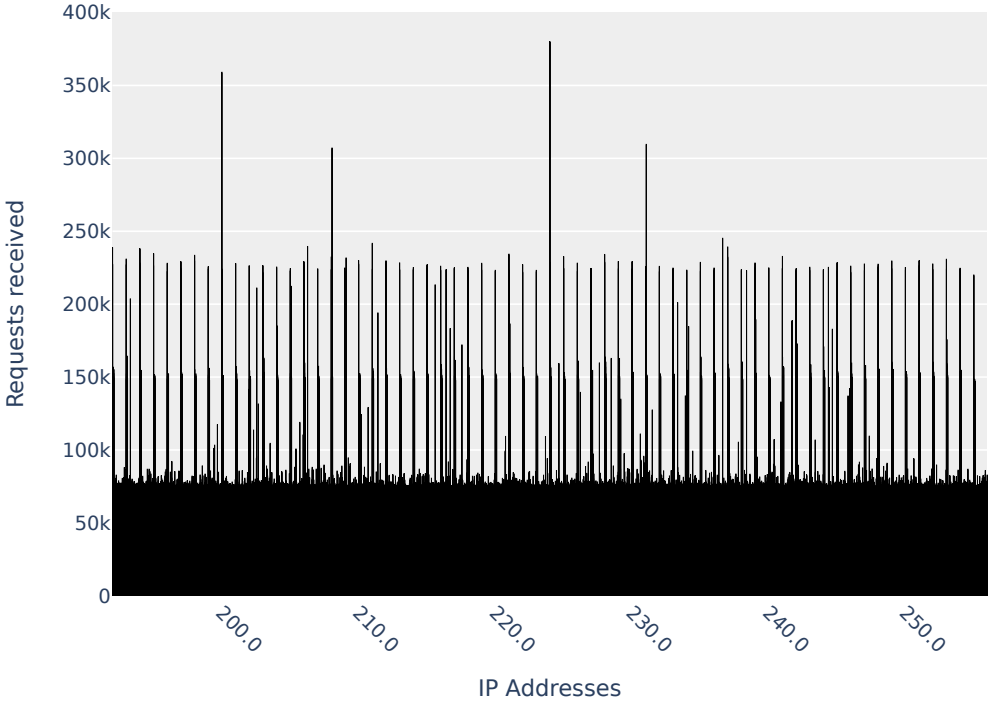

Figure 5.11: NICTER Sensor G, Requests per IPv4 (/21)

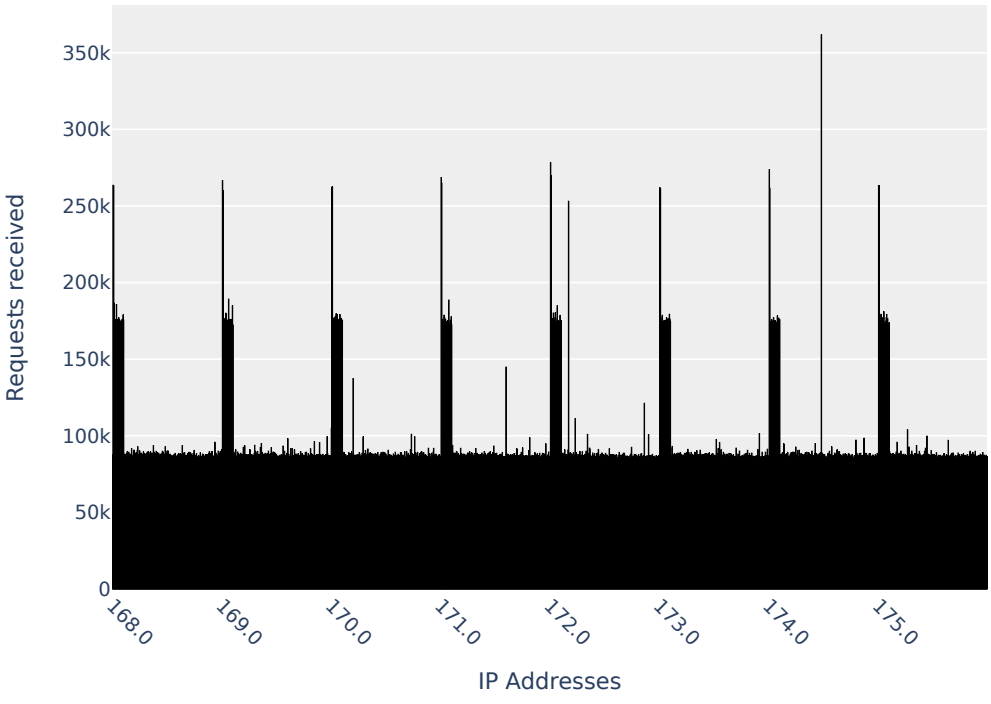

Figure 5.12: NICTER Sensor H, Requests per IPv4 (/21)

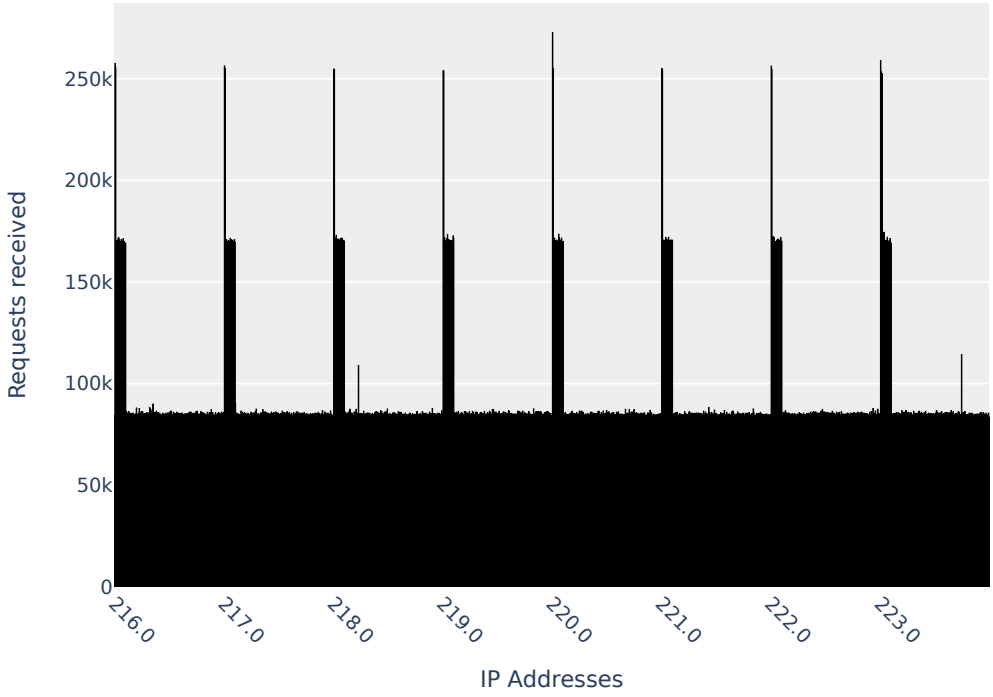

5.2.8 Sensor H

Sensor H, Figure 5.12 is the same size of subsection 5.2.7, they are very similiar, showcasing that the distribution of the sensors don't really differ much.

5.3 Analysis of IPv4 Address Space Reduction in Network Telescopes

As described in section 4.2, we explored the datasets and analyzed possible strategies to overcome problems related to the scarcity of IPv4 address space. Initially, we explored the distribution of the requests per destination IP in all the datasets we listed.

In fig. 5.13 we can visualize the distribution of requests received per IP address in both telescopes in a period of 31 days (one month). Here we depict Darknet-BR for the period of Dec/2023 and NICTER-E for Oct/2018. It's noteworthy that other NICTER sensors have shown a similar behaviour when compared with NICTER-E. In all datasets, the telnet scans (TCP/23) were the primary target; however, they were more prominent in 2018.

In Darknet-BR (fig. 5.13a), we received way more requests overall. It's probable

Figure 5.13: Distribution of received scans (TCP-SYN) per IP address in two different network telescopes. Both using a /19 block in one month.

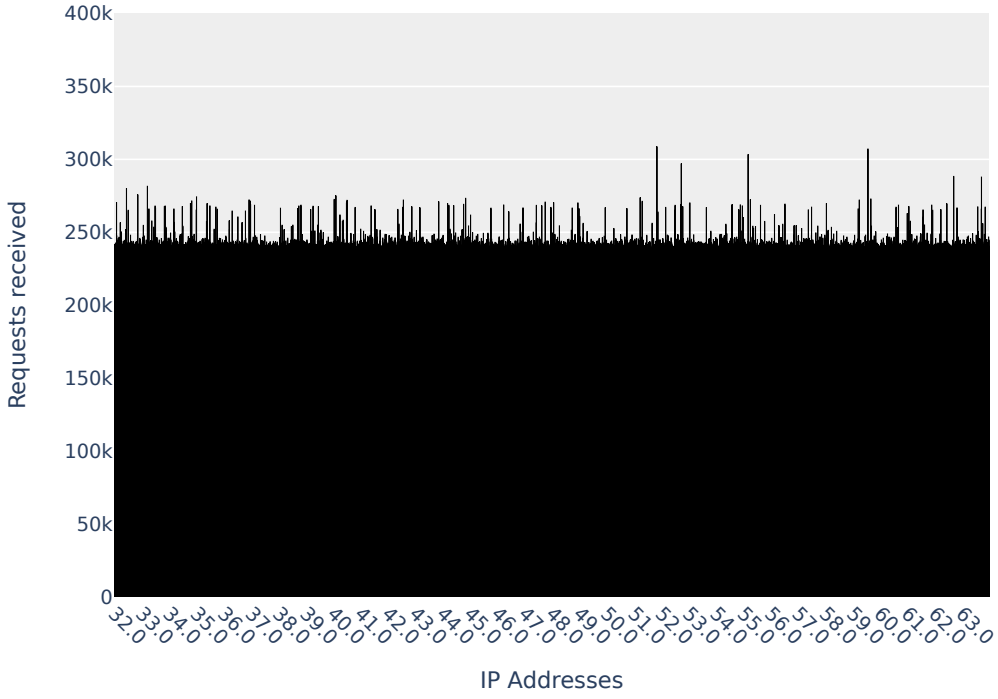

(a) Number of requests per IP, Darknet-BR.

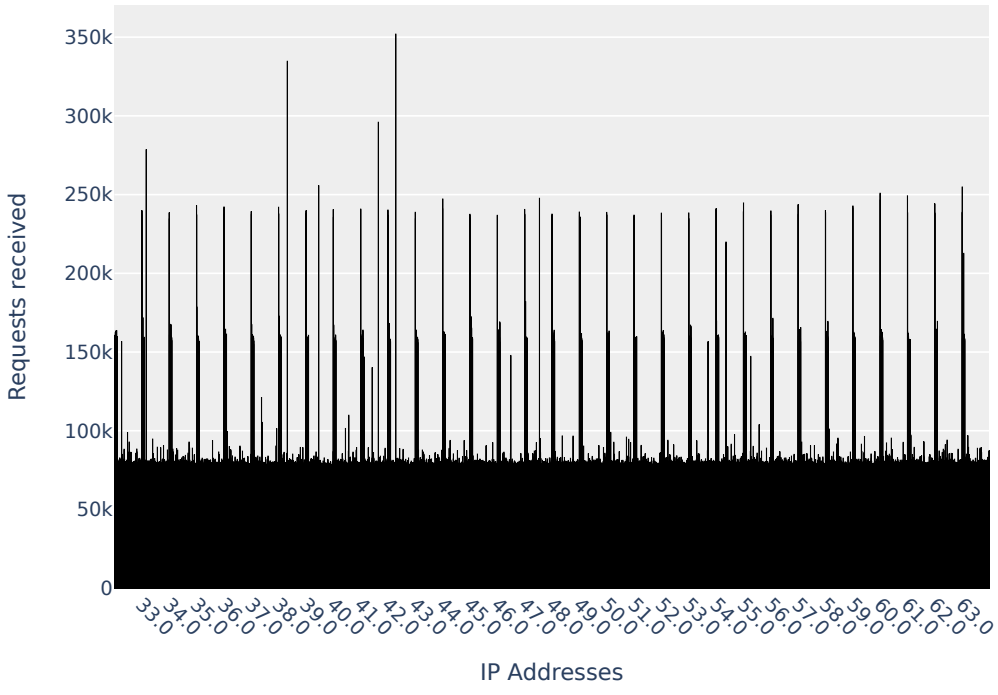

(b) Number of requests per IP, NICTER-E.

Figure 5.14: Darknet-BR /19 Percentage of unique sources received per IP used

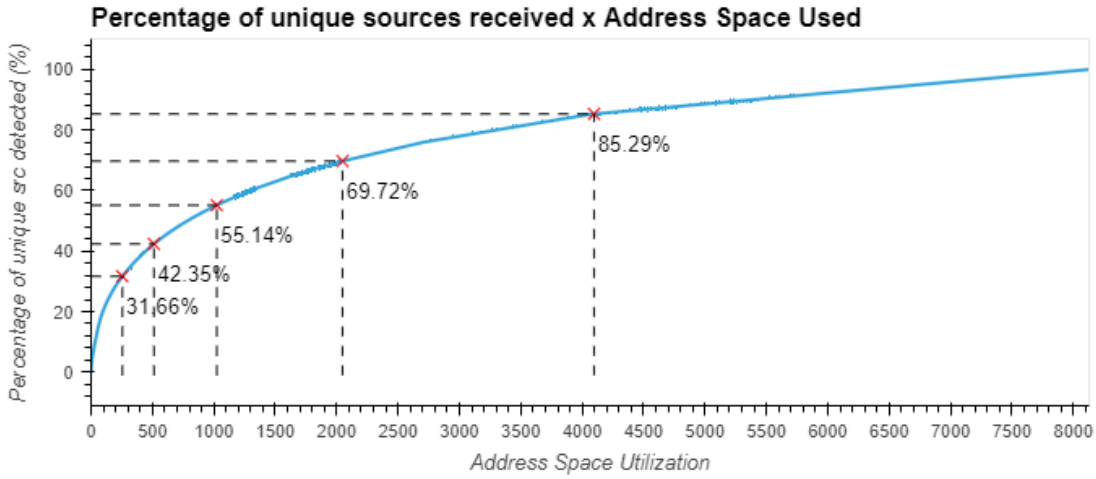

that the chronological factor is the main reason–more scans and malicious activities in 2023 than in 2018. Other point that worth to mention is the NICTER-E telescope show a discrepancy in the first 25 addresses of each /24 block as shown in section 5.2, indicating that most telnet scans target the first IPs in each /24 prefix, for example the IPs ending on .1 or .2. There were some cases in the literature (COOKE et al., 2004a), (CHINDIPHA; IRWIN; HERBERT, 2018) showing this increase of requests in lower IP addresses at the beginning of the /24 block, commonly allocated for gateways (*e.g.,* routers).

Applying the method described in section 4.2.3, we calculated the expected value of unique sources for all sizes of S for both sensors (Darknet-BR and NICTER-E). fig. 5.14 and Figure 5.15 show the percentage of the unique sources that would be visible when we project a reduction in the number of addresses being utilized by the telescope. The figure illustrate the projection of the percentage from the original telescope that we can reach (Y-axis) by the number of hosts needed (X-axis). This estimation allows us to assess how the reduction will impact the capture of unique sources by providing information on the origins already captured and the desired size of the new telescope.

When we compare the projected results for both telescopes, some similarities become evident. Both telescopes expect to capture more than 80% of unique sources when the address space is reduced by half, and more than 60% when reduced to a quarter. Our main observation here is that most of the identified scanned sources tend to scan several addresses in both telescopes.

Another observation is related to the incline of the curve in the two graphs. In this regard, NICTER-E shows slower growth in the beginning. This result is attributed to different attack methods, with NICTER-E registering most scans at lower addresses in each /24 (scans for routers). The takeaway here is that reducing the address space too

Figure 5.15: NICTER-E /19 Percentage of unique sources received per IP used

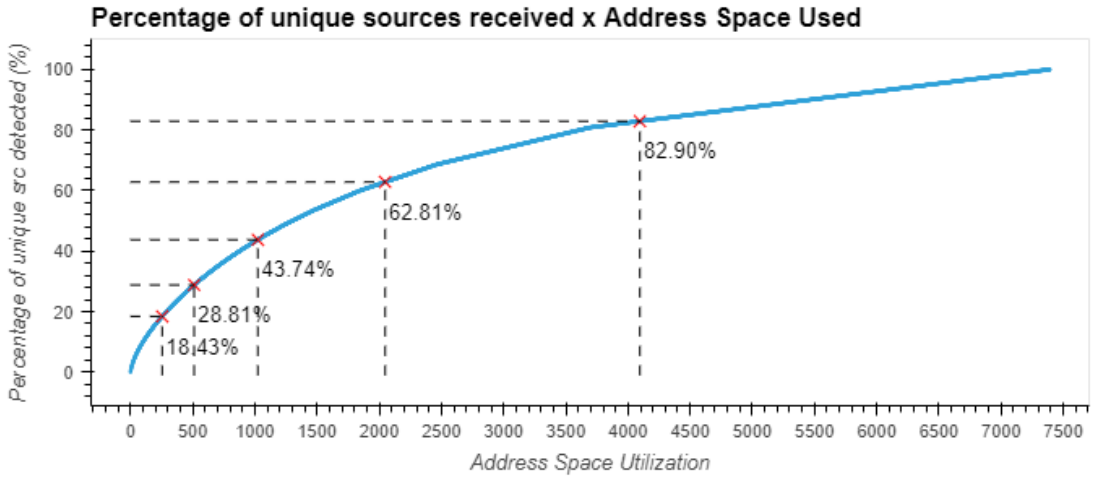

much may impact in our ability to detect certain types or methods of scans.

Additionally to the computation of the estimations used in section 4.2.3, we also relied on different sampling approaches to make those results more concrete. Here, we arranged the Darknet-BR dataset in 4 different subnets allocation schemas (e.g., one /20 or sets of /24), and gathered the real number of unique sources that would be perceived in each context, as described in section 4.2.4.

In table 5.4, we present the results of reducing the size of Darknet-BR by half, i.e., from 8k addresses to 4k addresses. The figure illustrates the impact on unique sources and the number of requests observed for each applied sampling method.

Table 5.3: Number of unique sources and requests seem by different methods.

Method	Unique Sources (%)	Number requests (%)
Total	100.00	100.00
Low /20	80.30	50.03
High /20	80.26	49.97
Even /24 allocation	80.26	50.01
Odd /24 allocation	80.39	49.99

Table 5.4: Number of unique sources and requests seem by different methods.

Method	Unique Sources (%)	Number requests (%)
Total	100.00	100.00
Low /20	74.73	50.01
High /20	74.46	49.98
Even /24 allocation	74.74	49.97
Odd /24 allocation	74.46	50.02

Our results from testing four different allocation methods over our Darknet-BR

dataset do not show a significant difference, less than 0.1% for identifying unique scan sources or requests. This occurs because of the uniform distribution present in the requests per IP addresses, meaning we did not find any particular set of IPs being more targeted than others.

We also examined the influence of the “border” and “middle” address space—/24 blocks at the beginning, end, or in the middle of the address space. Our objective was to determine if there is any difference between each /24. Although the first block does show more unique sources, it accounts only for 28.35%, while the worst /24 observed 28.11% of the total number of sources—a minimal difference of only 0.24%.

As to understand more about this behaviour, we also examined some individual /24 blocks in the beginning, end and in the middle of the /20 with the objective to assert that there are not a specific block that is significantly different from the others. Although the first block does show more unique sources, it just sees 28.35% while the worst /24 observed 28.11% of the total number of sources, a minimal difference of only 0.24%.

6 CONCLUSIONS

In this thesis, we propose to understand the usual range of address spaces utilized in network telescopes and to determine the impact of reducing address space on the quality of cyber threat detection. To achieve this, we address two research questions:

1. **What is the typical range of address spaces utilized in network telescopes ?**
2. **What is the impact of reducing address space in network telescopes on the quality of cyber threat detection?**

For the first research question, we studied the literature to gather information about all the networks telescopes that has been deployed in the last 23 years and taking in consideration the number of IPv4 they utilized. Additionally, we also collected information about the date of their creation and some historical data that identifies the telescope.

As a result of this research question, in section 5.1, we present as our contribution a list of 28 distinct initiatives that we have identified, of which 18 are still active today. Additionally, we draw a trend graph of these initiatives over the years, allowing us to illustrate how the depletion of IPv4 addresses is impacting network telescope initiatives and their sizes.

To address our second research question, we analyzed two network telescope datasets: the Darknet-BR from Brazil and several sensors from the NICTER network in Japan.

In section 5.2, we conducted a detailed analysis of the IP address distribution within the NICTER sensors. Our analysis focused on TCP-SYN scans due to limitations of the NICTER dataset. This analysis revealed a non-uniform distribution of scans across the IP address space, with at least 25 addresses appearing to be more targeted than others. Further examination of the datasets identified telnet scans for routers as the primary cause of this non-uniform distribution, a common attack in 2018.

From our initial analysis, we suspect that the bias towards lower IPs in NICTER could somehow impact the reduction in address space compared to the number of scans received in each network telescope.

Building on our previous experiment, we confidently selected NICTER-E as representative of all NICTER sensors. In section 5.3, we compared the NICTER-E sensor with Darknet-BR, as both share the same address space. Our comparison focused on requests per IP distribution, revealing that Darknet-BR receives more requests overall and

exhibits a more uniform distribution.

Lastly, we aimed to illustrate the impact of reducing address space on cyber threat detection. We proposed an estimated value formula in section 4.2.3 to assess the number of unique sources that would still be captured by the telescope despite the reduction. Additionally, we suggested applying sampling techniques to maximize the number of unique sources gathered by the sensors. Our findings, as shown in Figure 5.14, indicate that even with a reduction from a /19 to a /20 telescope in Darknet-BR, more than 80% of its total number of unique sources can still be detected. We replicated this method on NICTER-E, as depicted in Figure 5.15, and obtained similar results.

Furthermore, the addressing schema adopted in the reduction of Darknet-BR, such as splitting into multiple /24 or selecting one or another /20, has a minimal influence of less than 0.1%, as we depict in Table 5.4.

Our main contributions in this work are to offer a refreshed view on network telescope initiatives over the past 23 years and to estimate the impact of reducing a network telescope from a /19 to a /20 block size. Our tests demonstrate that these results are generalizable, and our method can be applied in other similar situations.

7 FUTURE WORKS

As future work, we intend to analyze other datasets to compare with our network telescope, that way it is possible to better understand how different telescopes behave while downsized. We would also like to observe the CAIDA dataset, which contains a very large network telescope, that way we could observe more throughly the process of reducing the size of a bigger telescope.

Furthermore, we will explore more protocols like UDP and look at more than just TCP-SYN data. That way, it will be possible to observe more patterns of attacks and exploits that utilize those protocols as their mean to the victims (ie. Reflection DDoS). There is also a possiblity to also utilize honeypots instead in order to have a more insightful view of how the attacks are conducted and to bring new insights to the field.

Last but not least, we intend to take a closer look at the locality factor of each network telescope as other reserches suggest (SORO et al., 2019) not all the telescopes are equal and their location and AS can have great impact on the data they collect. Another approach would be to study the hop distances that those requests normally are located relative to the sensors, that way it is possible to come with better solutions for attribution of more indirect attacks as Reflected DDoS.

REFERENCES

AFRINIC. **AFRINIC IPv4 Exhaustion statistics**. <https://stats.afrinic.net/ipv4/exhaustion/ipv4_available>. (Accessed on February 28, 2024).

AHMED, E. **Monitoring and analysis of internet traffic targeting unused address spaces**. Thesis (PhD) — Queensland University of Technology, 2010.

AHMED, E.; CLARK, A.; MOHAY, G. A novel sliding window based change detection algorithm for asymmetric traffic. In: **2008 IFIP International Conference on Network and Parallel Computing**. [S.l.: s.n.], 2008. p. 168–175.

ANTONAKAKIS, M. et al. Understanding the mirai botnet. In: **26th USENIX security symposium (USENIX Security 17)**. [S.l.: s.n.], 2017. p. 1093–1110.

APNIC. **Apic IPv4 exhaustion**. <https://www.apnic.net/manage-ip/ipv4-exhaustion/>. (Accessed on February 28, 2024).

BALKANLI, E.; ZINCIR-HEYWOOD, A. N. On the analysis of backscatter traffic. In: IEEE. **39th Annual IEEE Conference on Local Computer Networks Workshops**. [S.l.], 2014. p. 671–678.

CABANA, O. et al. Threat intelligence generation using network telescope data for industrial control systems. **IEEE Transactions on Information Forensics and Security**, v. 16, p. 3355–3370, 2021.

CAIDA. **Historical and Near-Real-Time UCSD Network Telescope Traffic Dataset**. 2024. <https://www.caida.org/catalog/datasets/telescope-near-real-time_dataset>. Accessed on February 28, 2024.

CENSYS Seach. <https://search.censys.io/>. (Accessed: February 28, 2024).

CHINDIPHA, S. D.; IRWIN, B.; HERBERT, A. Effectiveness of sampling a small sized network telescope in internet background radiation data collection. In: **Southern Africa Telecommunication Networks and Applications Conference (SATNAC)**. [S.l.: s.n.], 2018.

COOKE, E. et al. Toward understanding distributed blackhole placement. In: **Proceedings of the 2004 ACM workshop on Rapid malcode**. [S.l.: s.n.], 2004. p. 54–64.

COOKE, E. et al. The internet motion sensor: A distributed global scoped internet threat monitoring system. **Technical Report CSE-TR-491-04**, Citeseer, 2004.

CYMRU, T. **Team Cymru Darknet Project**. Accessed: February 28, 2024. Available from Internet: <https://www.team-cymru.com/>.

D'ANDRÉA, E. et al. Multi-label classification of hosts observed through a darknet. In: IEEE. **NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium**. [S.l.], 2023. p. 1–6.

EZE, D. T.; SPEAKMAN, D. L.; ONWUBIKO, D. C. **ECCWS 2020 19th European Conference on Cyber Warfare and Security**. [S.l.]: Academic Conferences and publishing limited, 2020. Google-Books-ID: lB4EEAAAQBAJ. ISBN 978-1-912764-62-4.

FACHKHA, C.; DEBBABI, M. Darknet as a source of cyber intelligence: Survey, taxonomy, and characterization. **IEEE Communications Surveys \& Tutorials**, Institute of Electrical and Electronics Engineers (IEEE), v. 18, p. 1197–1227, 2016. Available from Internet: <http://dx.doi.org/10.1109/comst.2015.2497690>.

Farsight Security. **Farsight Security, cyber security intelligence solutions**. Accessed: February 28, 2024. Available from Internet: <https://www.farsightsecurity.com/>.

FENG, Y. et al. A behavior-based method for detecting distributed scan attacks in darknets. **Journal of information processing**, Information Processing Society of Japan, v. 21, n. 3, p. 527–538, 2013.

GADHIA, F. et al. Comparative analysis of darknet traffic characteristics between darknet sensors. In: **2015 17th International Conference on Advanced Communication Technology (ICACT)**. [S.l.: s.n.], 2015. p. 59–64. ISSN: 1738-9445.

Google Cloud. **Google Cloud Virtual Private Cloud (VPC) Pricing**. 2023. <https://cloud.google.com/vpc/pricing-announce-external-ips?hl=pt-br>. Accessed: February 28, 2024.

GRIFFIOEN, H.; DOERR, C. Could you clean up the internet with a pit of tar? investigating tarpit feasibility on internet worms. In: **2023 IEEE Symposium on Security and Privacy (SP)**. [S.l.: s.n.], 2023. p. 2551–2565. ISSN: 2375-1207.

HAN, C. et al. Dark-tracer: Early detection framework for malware activity based on anomalous spatiotemporal patterns. **IEEE Access**, v. 10, p. 13038–13058, 2022.

HARDER, U. et al. Observing internet worm and virus attacks with a small network telescope. **Electronic Notes in Theoretical Computer Science**, Elsevier, v. 151, n. 3, p. 47–59, 2006.

HARROP, W.; ARMITAGE, G. Defining and evaluating greynets (sparse darknets). In: IEEE. **The IEEE Conference on Local Computer Networks 30th Anniversary (LCN'05) l**. [S.l.], 2005. p. 344–350.

HOUMZ, A. et al. Detecting the impact of software vulnerability on attacks: A case study of network telescope scans. v. 195, p. 103230, 2021. ISSN 1084-8045. Available from Internet: <https://www.sciencedirect.com/science/article/pii/S1084804521002290>.

HUIDES, A.; SANTHANAM, A.; LEHWESS, M. **Identify and optimize public IPv4 address usage on AWS | Networking & Content Delivery**. 2023. <https://aws.amazon.com/blogs/networking-and-content-delivery/identify-and-optimize-public-ipv4-address-usage-on-aws/>. Accessed: February 28, 2024.

IPv4 Global. **November 2023 Sales Report**. 2023. <https://ipv4.global/reports/november-2023-sales-report/>. Accessed: February 28, 2024.

IRWIN, B. V. W. A framework for the application of network telescope sensors in a global ip network. Rhodes University; Faculty of Science, Computer Science, 2011.

IUCC, T. **The IUCC/IDC Internet Telescope**. 2024. Acessed: February 28, 2024. Available from Internet: <https://nocvm.iucc.ac.il/research/telescope/>.

JONKER, M. et al. Millions of targets under attack: a macroscopic characterization of the dos ecosystem. In: **Proceedings of the 2017 Internet Measurement Conference**. New York, NY, USA: Association for Computing Machinery, 2017. (IMC '17), p. 100–113. ISBN 9781450351188. Available from Internet: <https://doi.org/10.1145/3131365.3131383>.

KALLITSIS, M. et al. Detecting and interpreting changes in scanning behavior in large network telescopes. **IEEE Transactions on Information Forensics and Security**, IEEE, v. 17, p. 3611–3625, 2022.

LACNIC. **Estadísticas de Asignación de LACNIC**. <https://www.lacnic.net/999/1/lacnic/>. (Accessed on February 28, 2024).

LAGRAA, S.; CHEN, Y.; FRANçOIS, J. Deep mining port scans from darknet. v. 29, n. 3, p. e2065, 2019. ISSN 1099-1190. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/nem.2065. Available from Internet: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nem.2065>.

MALÉCOT, E. L.; INOUE, D. The carna botnet through the lens of a network telescope. In: SPRINGER. **Foundations and Practice of Security: 6th International Symposium, FPS 2013, La Rochelle, France, October 21-22, 2013, Revised Selected Papers**. [S.l.], 2014. p. 426–441.

Merit Network. **Orion Network Telescope – Merit**. <https://www.merit.edu/initiatives/orion-network-telescope/>. Accessed: February 28, 2024.

MOORE, D. et al. Network telescopes: Technical report. 2004.

NICTER WEB. **NICTERWEB - Dark net observation | National Institute of Information and Communications Technology Cybersecurity Laboratory**. Accessed: February 28, 2024. Available from Internet: <https://www.nicter.jp/en>.

NIRANJANA, R.; KUMAR, V. A.; SHEEN, S. Darknet traffic analysis and classification using numerical AGM and mean shift clustering algorithm. v. 1, n. 1, p. 16, 2019. ISSN 2661-8907. Available from Internet: <https://doi.org/10.1007/s42979-019-0016-x>.

O'HARA, J. Cloud-based network telescope for internet background radiation collection. 2019.

PAULEY, E.; BARFORD, P.; MCDANIEL, P. DScope: A Cloud-Native internet telescope. In: **32nd USENIX Security Symposium (USENIX Security 23)**. Anaheim, CA: USENIX Association, 2023. p. 5989–6006. ISBN 978-1-939133-37-3. Available from Internet: <https://www.usenix.org/conference/usenixsecurity23/presentation/pauley>.

PEMBERTON, D.; KOMISARCZUK, P.; WELCH, I. Internet background radiation arrival density and network telescope sampling strategies. In: **2007 Australasian Telecommunication Networks and Applications Conference**. [S.l.: s.n.], 2007. p. 246–252.

RICHTER, P.; BERGER, A. Scanning the scanners: Sensing the internet from a massively distributed network telescope. In: **Proceedings of the Internet Measurement Conference**. [S.l.: s.n.], 2019. p. 144–157.

RICHTER, P.; BERGER, A. Scanning the Scanners: Sensing the Internet from a Massively Distributed Network Telescope. In: **Proceedings of ACM IMC 2019**. Amsterdam, Netherlands: [s.n.], 2019.

RIPE. **What is IPv4 Run Out?** <https://www.ripe.net/manage-ips-and-asns/ipv4/ipv4-run-out/>. (Accessed on February 28, 2024).

SHAIKH, F. et al. A machine learning model for classifying unsolicited iot devices by observing network telescopes. In: IEEE. **2018 14th International Wireless Communications & Mobile Computing Conference (IWCMC)**. [S.l.], 2018. p. 938–943.

SHODAN Search Engine. <https://www.shodan.io/>. (Accessed: February 28, 2024).

SORO, F. et al. Sensing the noise: Uncovering communities in darknet traffic. In: **2020 Mediterranean Communication and Computer Networking Conference (MedComNet)**. [S.l.: s.n.], 2020. p. 1–8.

SORO, F. et al. Are darknets all the same? on darknet visibility for security monitoring. In: **2019 IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN)**. [S.l.: s.n.], 2019. p. 1–6. ISSN: 1944-0375.

SWITCH. **SWITCH**. Accessed: February 28, 2024. Available from Internet: <https://www.switch.ch/>.

University of California San Diego - Center for Applied Internet Data Analysis. **Supporting Research and Development of Security Technologies through Network and Security Data Collection**. 2018. <https://apps.dtic.mil/sti/trecms/pdf/AD1054333.pdf>. (Accessed on 01/21/2024).

WAGNER, D. et al. How to operate a meta-telescope in your spare time. In: **Proceedings of the 2023 ACM on Internet Measurement Conference**. New York, NY, USA: Association for Computing Machinery, 2023. (IMC '23), p. 328–343. ISBN 9798400703829. Available from Internet: <https://doi.org/10.1145/3618257.3624831>.

WUSTROW, E. et al. Internet background radiation revisited. In: **Proceedings of the 10th ACM SIGCOMM conference on Internet measurement**. [S.l.: s.n.], 2010. p. 62–74.

ZAKROUM, M. et al. Self-supervised latent representations of network flows and application to darknet traffic classification. **IEEE Access**, IEEE, 2023.