Through a Smaller Lens: Revisiting Opportunistic Analysis using Network Telescopes

Bernhard Degen[1], Nils Kempen[2], kc claffy[3], Ricky K. P. Mok[3], Ralph Holz[2,1], Roland van Rijswijk-Deij[1], Raffaele Sommese[1], Mattijs Jonker[1]

[1] University of Twente, Enschede, The Netherlands
[2] University of Münster, Münster, Germany
[3] CAIDA/UC San Diego, La Jolla, CA, USA

Abstract. Unsolicited network traffic observed to the addresses monitored by a network telescope enables, among other things, tracking of Internet outages, botnets and DDoS attacks. We examine how a decrease in available address space affects what we can learn about the phenomena we study with telescopes. We conduct a targeted replication of a seminal study conducted 10 years ago. Since then, IPv4 scarcity and rising operational costs have placed increased pressure on operators to maximize use of their allocated space, which has resulted in a reduction of address space available to major telescopes. As a first step, we characterize traffic to three network telescopes that differ in size, spatial distribution, and prominence. We find that most address blocks within each telescope observe a similar number of source IP addresses, and that smaller telescopes offer higher visibility per monitored address. We also find that sources target the IPv4 address space pervasively, with 37.0% of them targeting a /16 block in each of the three telescopes within an hour. As a case study, we examine the sensitivity of randomly-spoofed DoS attack inference to the size of the address space under observation and find that larger telescopes detect many attacks missed by smaller ones, although smaller telescopes observe disproportionately many relative to their address space size. Our study provides a framework to quantify the effects of reduced telescope address space and outlines future directions for telescope research.

1 Introduction

Network telescopes (also called *darknets*) are routed but unused address ranges. The unsolicited traffic they receive, referred to as Internet background radiation (IBR), has proved instrumental in identifying scans that probe for vulnerabilities [25,2,36,13,14,1], inferring Internet-wide randomly-spoofed denial-of-service (RSDoS) attack activity [24,21], and detecting Internet outages [11,3]. Operators can also use them to infer scanning activity targeting their networks [16,6].

In contrast to honeypots, which are designed to respond to probes, network telescopes are intentionally silent. As they consist of unused addresses, they

attract activity from misconfigured hosts, backscatter resulting from RSDoS attacks, and scans that indiscriminately target a sufficiently large part of the IPv4 address space.

In recent years, we have seen a confluence of issues that threaten to affect the use of telescopes as a tool to study Internet phenomena. First, a growing body of work shows that scanners target different parts of the Internet differently [29,20].

Second, the scarcity of IPv4 address blocks [27] has put pressure on operators to maximize use of their allocated space. In the last 10 years, operators of two major network telescopes (`UCSD-NT` and `MERIT-NT`) have seen significant reductions of their address space. Such reductions may limit the effectiveness in observing Internet phenomena such as outages and security threats.

Third, over the past two decades, the volume of traffic captured by large telescopes has grown considerably [15], placing increasing demands on IBR archival and analysis infrastructure. These challenges hinder researchers' ability to extract meaningful insights from IBR, providing additional motivation to investigate how telescope size and placement affect visibility of Internet-wide activity.

These challenges make it imperative to understand 1) how representative IBR observed in a network telescope is of IBR destined to the entire address space, and 2) how a telescope's capacity to capture global events depends on the size of its address space.

This paper revisits Benson et al. [4], a seminal work published in 2015 on analyzing IBR, to assess the validity of the findings in today's traffic dynamics and telescope deployments. Specifically, we address the following research questions: How many sources are observed over time (§6.1)? How frequently do sources contact telescopes (§6.2)? How large are the observed autonomous systems (ASes), based on their advertised address space (§6.3)? How does the visibility of telescopes depend on their sizes (§6.4)? To what extent do sources target[4] specific parts of a telescope's address space (§6.5)? To what extent are sources observed across telescopes (§6.6) and how does IP/port concentration differ (§6.7)? In addition, we perform a case study where we examine the sensitivity of RSDoS inference to the size of the address space (§6.8). We use data from the same two large, but downsized network telescopes (`UCSD-NT` and `MERIT-NT`) and from one additional, smaller telescope in a topologically and geographically different location (`SURF-NT`), and re-examine the question of how telescope size and position influence visibility. Our findings are largely consistent with those of the original study. While we observe higher data volumes and packet rates, network telescopes continue to provide a valuable lens for Internet-wide analysis. Our contributions are as follows:

1. We provide an empirical comparison of three network telescopes varying in size, distribution, and prominence, identifying differences in the volume and patterns of sources targeting them.

[4] We use "target" to indicate that traffic is directed toward the telescope and do not imply any intent by the source.

2. We find that many sources target the IPv4 address space indiscriminately: 37.0% of source IP addresses targeted a /16 block across three network telescopes within one hour.
3. We conduct a case study on the inference of RSDoS attacks and find that while larger address blocks captured more targets, a smaller telescope identified a disproportionally high share of them.
4. We discuss challenges faced by telescope operators and the efficacy of alternative deployment strategies in overcoming these.

2 Background

2.1 Network Telescopes

Network telescopes are collections of routed but unused address space that passively collect traffic. Since they do not host any services, they capture Internet background radiation (IBR): persistent, unsolicited traffic. This includes errant packets from misconfigured hosts, backscatter from RSDoS attacks, and scanning activity. This traffic provides academia and industry with a valuable resource for collecting cyber threat intelligence (CTI). Telescopes can be *centralized*, with all monitored addresses in a single block, or *distributed*, spanning multiple non-contiguous blocks, possibly in diverse topological or geographical regions. These deployment choices influence the type and volume of traffic they attract [29,20].

With increased global scanning and attack activity, the volume of IBR has surged. At the same time, IPv4 scarcity limits the address space available for telescopes. Despite reductions in telescope address space, between 2002 and 2022 IBR captured by `UCSD-NT` has increased by three orders of magnitude [15], significantly raising the challenges of storing, processing, and analyzing telescope data.

Apart from IPv4, telescopes can also be deployed on unused IPv6 address space. However, exhaustively scanning the entire IPv6 address space is infeasible. Therefore, these telescopes attract a different class of scanners; those that use public data sources to select prefixes to scan [33].

2.2 RSDoS Inference

Network telescopes offer a valuable tool for inferring ongoing randomly-spoofed denial-of-service (RSDoS) attacks. In an RSDoS attack, an attacker sends traffic with randomly spoofed source addresses to its target, aiming to overwhelm it. Spoofing the source addresses can enable the attacker to circumvent traffic filters or thwart intervention. This type of attack is visible to telescopes via *backscatter*: reply packets from the target to a spoofed address. The `UCSD-NT` covers approximately 0.254% of the IPv4 address space, and assuming source addresses are spoofed uniformly at random, there is an equal likelihood of 0.254% that a reply packet will be directed to an address within the telescope's range. The probability of observing at least one backscatter packet is

$$1 - \left(1 - \frac{n}{2^{32}}\right)^m$$

where n is the number of addresses monitored by the telescope and m the number of packets sent by the attacker. For example, a modest attack of 1,000 packets per second (pps) lasting one minute would be observed in a /16 telescope with a probability of approximately 60% and in `UCSD-NT`'s entire /9 and /10 address blocks with nearly 100% probability.

By grouping reply packets by target IP address and recording when the first packet was observed, it is possible to estimate the duration and intensity of an attack [24]. Applying thresholds on the number of packets, rate, and duration can remove noise and small attacks with negligible impact.

3 Related Work

We group prior research along three dimensions: visibility across telescopes, scanner strategies, and efforts to collect IBR using fewer addresses. We build on these studies by revisiting a 10-year-old study under current conditions.

3.1 Visibility across Network Telescopes

Early works [25,2,36] characterized IBR in several /8 prefixes, noting differences in the ports targeted and the locality of the traffic. Several studies have shown that the geographical location of observation influences the IBR it received [32,17]. Like these studies, we evaluate visibility from multiple vantage points. In light of ongoing address space contraction, we further analyze intra-telescope visibility to assess the effectiveness of smaller telescopes.

3.2 Scanning Strategies

The characteristics of IBR can vary across different parts of the address space. Although the introduction of Internet-wide scanning tools has lowered the barrier to probing the entire IPv4 address space, scanners may still selectively target certain portions of it. One study [14] found that between January 1, 2013 and May 1, 2014, 68% of scans targeted at least 10% of the IPv4 address space, and another [29] analyzed server firewall logs from a content distribution network (CDN) in November 2014 and found that 30% of the scans were localized. Izhikevich et al. [20] studied how scanners discriminate between `MERIT-NT`, two educational networks, and honeypots hosted in five cloud environments during July 1–7, 2021, and found that networks with legitimate services attract more scanners. While these studies focused on scanner strategies and their biases toward certain network types, our study centers on how targeting behavior varies across different network telescopes.

3.3 Reducing Address Space Footprint

Traditional large-scale network telescopes passively monitor vast swaths of unused address space. Several alternative approaches have been proposed to reduce the address space footprint while maintaining or improving visibility. One

approach is to employ partially unused network resources, such as *greynet* addresses [16] or ports [29]. A related approach is to deploy sensors in cloud networks [5,20,26]. While these may draw scanners that do not appear in telescopes, the traffic a virtual machine (VM) receives can be influenced by the prior usage of its address. It is also possible to monitor IBR without owning the address space targeted, using Internet exchange points (IXPs) as vantage points [35]. However, this method's visibility is limited to traffic traversing the IXP and may inadvertently capture production traffic.

Similar to our work, several studies have explored the effect of downsizing existing telescopes. Chindipha et al. [9] used error metrics to compare /27–/30 samples to their covering /24 telescope in February 2018 and March 2019, demonstrating the viability of smaller sensors. Camargo et al. [6] analyzed /19 telescopes in Brazil (December 2023) and Japan (October 2018), showing that halving each still captured 80% of source addresses. It is unclear whether these studies filtered spoofed packets, which can heavily skew results (§5.2).

3.4 Leveraging Internet Background Radiation for Opportunistic Network Analysis

Benson et al. [4] compared the utility of IBR for opportunistic network analysis in two network telescopes, `UCSD-NT` and `MERIT-NT`. They analyzed `UCSD-NT` for two one-month periods — one in 2012 and the other in 2013 — and `MERIT-NT` for the coinciding one-month period in 2013. Their work quantified sources, analyzed IBR components, and characterized temporal IBR patterns. They found an overlap of 84% in observed source /24 blocks in both telescopes, but noted substantial variation across `UCSD-NT` /16 blocks. Ten years later, the `UCSD-NT` and `MERIT-NT` have shrunk by 34% and 96%, respectively. We revisit the utility of these network telescopes for Internet-wide analysis using updated data as part of a targeted replication of their work, and we include one additional, smaller network telescope (§4). Motivated by the persistent threat of distributed denial-of-service (DDoS) attacks and the use of telescopes to study them [31,18], we benchmark RSDoS inference as a case study.

4 Datasets

We utilize three network telescopes and two additional datasets.

`UCSD-NT` A major network telescope composed of a /9 and an adjacent /10 subnet, interspersed with assigned subnets. 13.2% of the addresses belong to these assigned subnets and are therefore excluded from our dataset, leaving 10.9 M monitored addresses.

`MERIT-NT` A network telescope composed of various /14–/23 subnets, some of which are contiguous, totaling 476 k addresses.

`SURF-NT` A network telescope composed of one /16 and three /24 non-contiguous subnets, totaling 66.3 k addresses.

Table 1: Starting times (UTC) and packet counts of random one-hour samples.

Sample		Packet count		
Date	Time	UCSD-NT	MERIT-NT	SURF-NT
2025-04-05	03:00	3,545,835,110	209,260,789	25,210,194
2025-04-06	11:00	3,409,739,023	200,977,859	29,848,213
2025-04-07	16:00	3,314,010,076	179,474,838	21,243,581
2025-04-08	22:00	4,434,768,231	176,348,784	19,603,466
2025-04-09	05:00	4,030,952,660	187,660,477	22,591,846
2025-04-10	14:00	3,457,389,678	202,179,146	26,560,589
2025-04-11	02:00	3,356,122,335	211,108,787	28,547,691

BGP RIB We use a Border Gateway Protocol (BGP) routing information base (RIB) to resolve the AS originating an IP address and to count advertised prefixes by their size. Specifically, we use the RIB from collector `route-views2.routeviews.org` [30] captured on April 5, 2025 at 00:00 UTC, corresponding to the start of our observation period.

RIR allocation statistics We use allocation statistics published by each regional Internet registry (RIR) on April 5, 2025 [34] to count allocations by prefix size.

The telescope datasets span seven days, starting on April 5 at midnight UTC. The total gzip-compressed size of the pcaps is 20.4 TB, 2.25 TB, and 121 GB for `UCSD-NT`, `MERIT-NT`, and `SURF-NT`, respectively. While our observation period is shorter and the telescopes are smaller than in the original study [4], we process 4–5.1× and 1.5× more data for `UCSD-NT` and `MERIT-NT`, respectively. Due to gaps in the availability of `UCSD-NT` pcap files within the observation period and constraints on processing and storage resources, we use seven random (but consistent throughout the study) 60-minute samples, one from each day. This sampling strategy was designed to maximize the amount of data we could process within our budget and available cluster resources. Inevitably, this choice sacrifices some visibility into longer-term trends, underscoring the trade-offs that researchers must make when analyzing large-scale network telescope data. For experiments where capturing the complete timespan is required (§6.1–6.3), we use flow data containing the necessary information for the full seven-day period[5].

Table 1 lists the times and packet counts for each random sample. The packet count per destination address is largely stable across both samples and telescopes. To capture the average pattern over time, we report the *mean* across time samples. To reflect the typical case while minimizing the influence of outliers (e.g., due to filtered telescope addresses), we report the *median* across address blocks. Although our results show consistent trends across samples, supporting

[5] `UCSD-NT` captures were unavailable during 2025-04-11 18:00–23:59. Where this gap affects our analysis, we note it explicitly.

the representativeness of our sampling method, we cannot rule out the possibility of anomalies or transient behaviors within the observation period.

5 Methodology

In this section we outline our methodology. We first explain how we preprocess packet captures (§5.1). Then, we explain how we correct for spoofed traffic targeting the network telescope itself, as such traffic can heavily skew the results of any analysis performed on captured traffic (§5.2).

5.1 Preprocessing Packet Traces

Each of the three telescopes provides raw packet traces in pcap format. `MERIT-NT` data is also available in the form of aggregated *events* [22], such as periods of scanning and backscatter. `UCSD-NT` data is additionally distributed as flows in the *FlowTuple* [7] format. Our study does not require payloads, as payload-based traffic classification is non-trivial and lies beyond its scope.

To obtain one unified data representation for all telescopes, we used Corsaro 3 [10] to convert raw traces to FlowTuples while preserving the destination IP address. For `UCSD-NT` analyses spanning the entire observation period, we used the provided FlowTuples with flows aggregated by destination /16 subnet to overcome storage and processing constraints. This granularity matches the requirements of our analyses.

5.2 Filtering Spoofed Source Addresses

For our analysis, we are not interested in traffic that is spoofed directly toward a network telescope. Unlike RSDoS attack traffic — which is spoofed toward a target, with the target's replies potentially reaching a telescope — this spoofed traffic is directed to a telescope itself. In this section, we describe how we mitigated the impact of such spoofed traffic on our results.

While it may be impossible to accurately detect spoofed traffic without conducting active measurements, previous work [12,10] empirically derived a set of heuristics to identify spoofed packets passively. These heuristics were based on commonalities in spoofed traffic captured in 2012. We tested them but found them unsuitable for our purposes for several reasons. First, some heuristics filtered out valuable IBR components because they exclude packets sent to UDP port 80, which presently is one of the ports assigned to HTTP over QUIC [19]. Second, we observed that the heuristics did not reliably identify spoofed packets during short-lived spoofing events. Although they flagged some sources above the baseline as spoofed, noticeable spikes in source addresses remained after filtering, indicating that the heuristics did not account for the full deviations. Given these limitations, we decided not to apply these heuristics.

Instead, we devised a statistics-based method to remove surges of spoofed traffic from our datasets, focusing only on short-lived, high-rate events that disproportionately skew our results. By targeting only these bursts, we minimize the

risk of excluding legitimate data throughout the observation period. To identify such events, we manually inspected the time series of unique source addresses and empirically derived a threshold of twice the mean number of sources. We computed the mean number of unique source addresses per 5-minute interval, the FlowTuple aggregation interval, and identified two high-rate spoofing events in `UCSD-NT`, along with several smaller ones. These events peaked at 79.0× and 18.0× the mean number of sources observed in each (excluding the events themselves). Both events occurred outside the one-hour samples listed in Table 1 and thus could only affect analyses based on data from the entire observation period (§6.1–6.3). For each event, we derived a signature based on the 5-minute time intervals in which it occurred, along with the targeted telescope addresses, ports, and protocols. After excluding these two high-rate spoofing events, only modest spikes (less than twice the mean number of sources) remained. While any residual spoofed traffic may slightly affect our results, removing these smaller deviations could come at the cost of removing legitimate IBR. We also inspected `MERIT-NT` and `SURF-NT` time series, but these datasets did not contain any intervals where the source address count exceeded the threshold; hence, we did not exclude any events. A possible explanation is that these telescopes are much smaller and therefore less likely to attract spoofed traffic than `UCSD-NT`.

More generally, there is an inherent risk that some persistent, low-rate spoofing remains undetected. Such residual spoofed traffic could lead to an overestimation of visibility in terms of sources observed, although this effect is limited relative to the excluded high-rate events.

6 Results

We first study how analyzing network telescopes over different time intervals impacts visibility, i.e., how many IBR sources it observes (§6.1). This helps us understand the trade-off between expending computational resources on longer analyses against the utility of the outcomes. Next, we analyze how frequently sources contact the three telescopes in our study (§6.2), which informs the predictability of IBR. We then analyze the source ASes observed by each telescope (§6.3) to estimate the coverage of the routed address space. We also investigate how visibility varies with the size (§6.4) and position (§6.5) of telescope address blocks, providing insight into the marginal gain in visibility and the similarity between blocks. Following that, we assess the locality of IBR to the blocks under observation (§6.6). This provides insight into how distributed traffic sources are over the IPv4 address space. After that, we analyze how probing patterns differ by telescope (§6.7) to better understand what draws traffic to specific regions of the address space. Finally, we conclude this section with a case study on RSDoS inference across blocks of varying sizes (§6.8), demonstrating the practical utility of different telescope configurations.

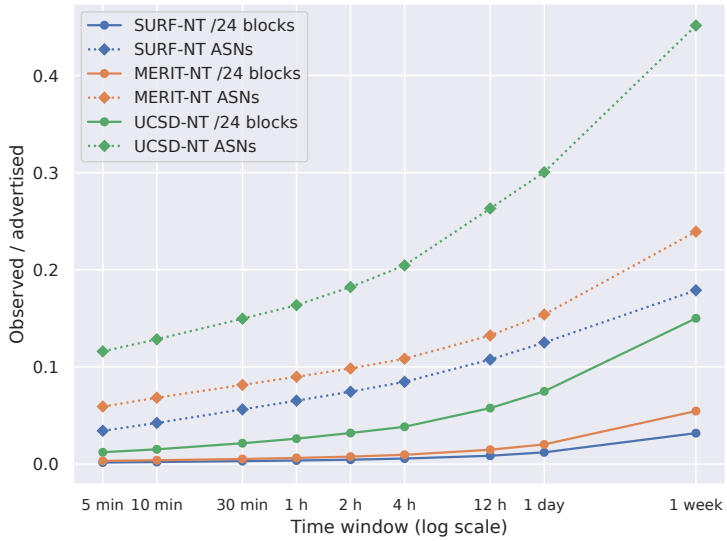

Fig. 1: Median autonomous system numbers (ASNs) and /24 blocks observed across time windows during the observation period. The observation counts increase for longer windows. The fractions are relative to the total number of advertised ASNs and /24 blocks. The faster-than-logarithmic growth suggests that longer observation periods capture more source diversity.

6.1 Effect of Observation Duration

To quantify the sources each network telescope attracts over time at different levels of granularity, we used two metrics: the number of observed source /24 blocks and ASNs. For each telescope, we calculated the median of these values across varying time windows. As shown in Figure 1, extending the observation period resulted in an increase in both /24 blocks and ASNs. Consistent with the study by Benson et al. [4], the number of sources observed grew sublinearly as we lengthened the time windows, due to sources repeatedly contacting a telescopes. As they also observed, however, the curves continued to rise near the longest window of one week, suggesting that even longer observation periods would yield additional sources. This trend was consistent across telescopes. Observations during one-week windows were between 3.51× (`MERIT-NT`) and 4.23× (`SURF-NT`) those of five-minute windows for ASNs, and between 9.91× (`UCSD-NT`) and 15.5× (`SURF-NT`) for /24 blocks. To illustrate, achieving a coverage of at least 10% of routed ASNs would necessitate observing traffic for at least 5 minutes in `UCSD-NT`, 2 hours in `MERIT-NT`, and 12 hours in `SURF-NT`. This shows that telescopes with smaller address spaces can achieve high source visibility over time.

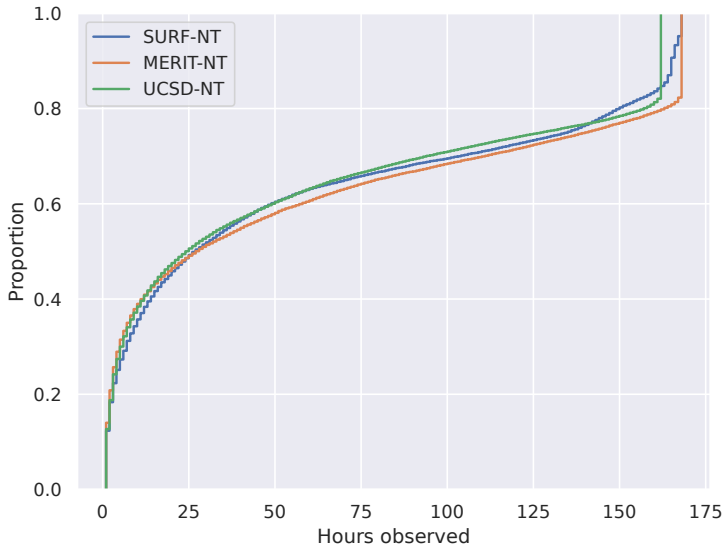

Fig. 2: ECDF of one-hour bins that ASNs were observed in. Approximately 18% of ASes contacted a telescopes every hour or more frequently, supporting longitudinal studies.

6.2 Contact Frequency of IBR Sources

The incessant nature of IBR has enabled longitudinal studies, such as network outage detection [11,3]. The accuracy of such inferences depends on the stability and frequency with which remote systems send traffic to network telescopes. To measure this, we counted the number of hours sources were observed and calculated the median time between these observations.

The FlowTuples used in this section contain packet counts per five-minute flow. To approximate packet timing within each flow, we assumed a uniform distribution of packet arrival times. The steps at multiples and simple fractions of 300 seconds in Figure 3 arise from this assumption.

Autonomous System Stability Repeated ASes observations enable longitudinal studies. Figure 2 shows the eCDF of the number of one-hour bins in which ASNs were observed over the one-week observation period. `UCSD-NT` was missing data for six hours; this gap is reflected in the lower maximum value. The tail of the distributions indicates that approximately 18% of ASes sent traffic at least hourly, making these ASes suitable for analyses that require regular observations. In contrast, /24 blocks are less reliable for longitudinal analysis, since we found only 2% of these active every hour. These findings align with the original study [4] and indicate that only a few sources exhibit sufficient stability to support longitudinal analysis requiring hourly observations. Alternatively,

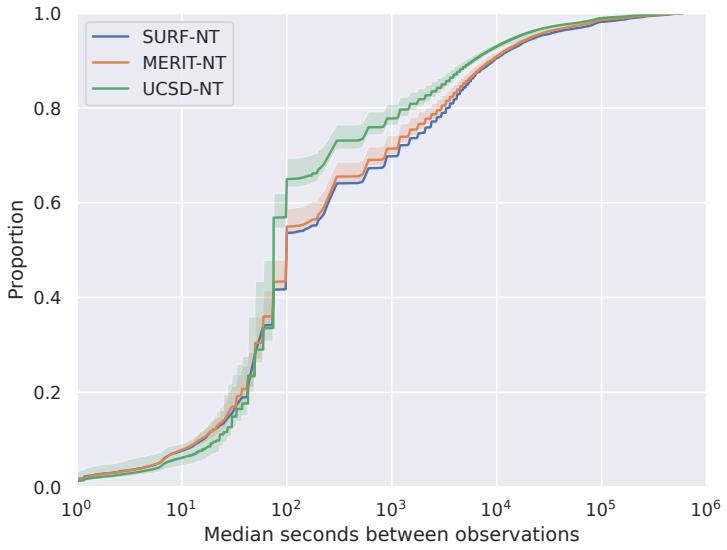

Fig. 3: ECDF of the median IAT per source /24 block in each /16 telescope block. The lines represent the median value across blocks, and the shaded regions indicate the minimum and maximum IATs. /24 blocks sent traffic to `UCSD-NT` at a higher rate than to `SURF-NT` and `MERIT-NT`. Nearly a third of /24 sources sent traffic with median IATs of one minute, supporting analysis that requires short intervals.

aggregation can be performed on longer time bins to trade precision for broader coverage.

Packet Rate Inferences may additionally benefit from short intervals between observations. Figure 3 shows the eCDF of the median inter-arrival times (IATs) of source /24 blocks. To account for the size difference of the three telescopes in our study, we calculated the IATs of packets destined to each /16 block fully covered by the respective telescope throughout the observation period (83 blocks for `UCSD-NT`, 7 for `MERIT-NT`, and 1 for `SURF-NT`). Between 27.6% and 30.2% of /24 blocks and 35.7% and 37.5% of ASes sent traffic with a median IAT of one minute or less. Although the original study [4] analyzed the entire `UCSD-NT`, which was 256× larger than the /16s we used, we consistently observed higher per-source packet rates.

In our datasets, `UCSD-NT`/16 blocks received traffic at higher rates than the /16s of the other two telescopes did. To analyze this difference, we examined the packet size distributions, since smaller packets can be transmitted at higher rates. While `UCSD-NT` observed approximately 10% more of the smallest packets (smaller than 44 bytes) relative to the other telescopes, we found that it received approximately 15% *fewer* medium-sized packets (44 to 60 bytes), which may

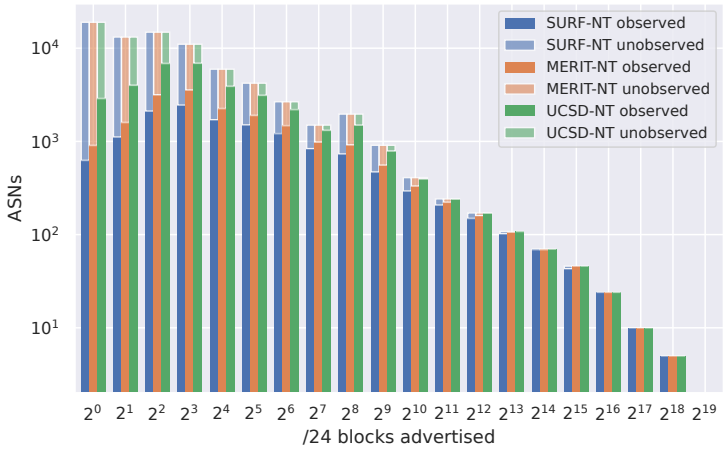

Fig. 4: ASes observed transmitting to network telescopes over the full observation period, grouped by number of originated IP address blocks, expressed in /24 equivalents. The solid and faded bars together indicate the number of ASes in that size class in the global routing table. `UCSD-NT` saw the most ASNs and `SURF-NT` the least. ASes originating many IP addresses were more likely to send traffic to the telescopes.

underlie the difference in IATs. At the AS granularity, telescopes exhibited more consistency with 80% of the ASes having a median IAT of at most 50, 45, and 45 minutes across `SURF-NT`, `MERIT-NT`, and `UCSD-NT` blocks, respectively.

In conclusion, only a small fraction of sources consistently sent traffic to a telescope every hour, whereas most sources exhibited median IATs on the order of minutes. This pattern suggests that IBR sources typically generate traffic in bursts, and that longitudinal inferences should be drawn with caution, especially at /24 or finer source granularity.

6.3 Autonomous System Visibility

To better understand how much of the routed address space is observable in each network telescope, we calculated the number of source ASNs, grouped by the size of their originated address space. Figure 4 shows the number of ASNs sending traffic to the three telescopes during one week. The size of the originated address space is expressed in powers of two of /24 blocks, rounded down. For example, an AS originating two /15s and another AS originating a /14 and a non-overlapping /15 are both counted as originating 2^{10} /24 blocks.

Differences among telescopes were most evident for ASes that originate relatively few addresses. During the one-week observation period, `UCSD-NT` captured traffic from at least half of the ASes originating address space equivalent to a /21 prefix or larger. `MERIT-NT` and `SURF-NT` observed many fewer small ASes;

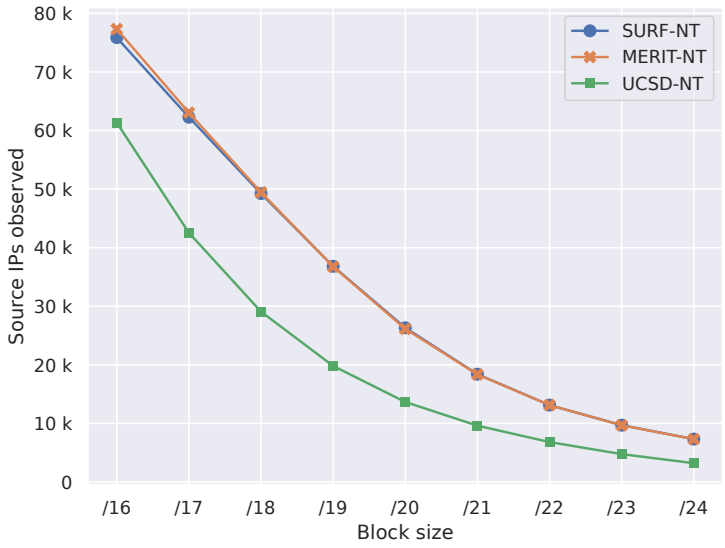

Fig. 5: Median IP addresses observed per hour across all /16–/24 network telescope blocks. The observation counts decrease for larger blocks. The rate of decrease slows with increasing x (halvings of the address space), indicating diminishing returns.

each captured at least half of the ASes originating address space equivalent to a /15 or larger. All network telescopes received traffic from the largest ASes, those originating address space equivalent to a /8 or larger.

Consistent with the original study [4], we found a relatively low visibility of ASes announcing an equivalent of a /16 (2^8 /24 blocks). This effect may stem from a concentration of legacy address space in this size class, which tends to exhibit lower active use. To investigate this further, we analyzed RIR allocation statistics and found that across all RIRs, there were 3.13× more /16 than /17 allocations, and 5.71× more /16 than /15 allocations. This may account for the increase of ASes in the /16 size class that we observed. Moreover, the RIR statistics reported a higher proportion of legacy address space among /16 allocations. Legacy allocations predate the existence of the RIRs and are often address blocks allocated to academic institutions and other early adopters of the Internet. This means that these allocations are likely generous in comparison to more recent allocations and as such see less dense utilization [28], providing an explanation for the relatively low visibility in the /16 size class.

6.4 Effect of Telescope Size

Existing network telescopes vary greatly in size [6]. Larger telescopes potentially capture a broader set of network activity, enabling more comprehensive insights.

To estimate the effect of size, we counted the source IP addresses observed in each telescope for all telescope block sizes from /16 to /24, using daily one-hour samples. For consistency, we only included /16 telescope blocks fully covered by the respective telescope. Given that the number of data points varies across block sizes and telescopes, we present only the median value for each block size in our analysis.

Figure 5 shows the results of this experiment. While the address space of each telescope is halved at each step, the number of sources did not decline as quickly. As also observed by Benson et al. in 2015 [4], we found that the number of sources scales as a power-law with the size of the address space.

To normalize for block size, we report the ratio of source addresses observed to destination addresses within the block. For `MERIT-NT` and `SURF-NT`, a destination address in a /24 block observed a median of 28.5 and 28.6 source addresses per hour, respectively, whereas an address in a /16 block saw only a median of 1.2 source addresses per hour for both telescopes. `UCSD-NT` destination addresses observed fewer sources than the other two telescopes: addresses in /24 and /16 blocks saw a median of 12.5 and 0.9 source addresses per hour, respectively. These findings indicate that increasing the monitored address space leads to diminishing marginal returns.

Additionally, we found that `UCSD-NT` observed fewer sources than the other two telescopes across all block sizes. The median number of source addresses observed per destination address was at least 19.9% lower compared to the same block sizes of the other two telescopes, suggesting that some IBR sources do not target the `UCSD-NT` address space.

6.5 Effect of Position in Address Space

To isolate the effect of position in the global IPv4 address space on observed sources, we divided the first `MERIT-NT` /16 block and the sole `SURF-NT` /16 block address space into smaller /24 blocks, counted the sources addresses observed in each /24 block, and calculated the mean over the seven daily one-hour samples. We repeated this procedure for `UCSD-NT` by dividing its covering /8 into /16 blocks. Several blocks were in assigned and therefore excluded from our dataset, resulting in fewer observed packets. Figure 6 shows the three resulting Hilbert curves. The number of source addresses was rather consistent across blocks, with only a few exceptions.

Figures 6a and 6b show /24 blocks within the first /16s of `MERIT-NT` and `SURF-NT`. It shows that, for both telescopes, the –.–.255/24 block observed fewer source addresses than the other /24 blocks. In addition to /24 blocks shown in the figure, we examined all /24 blocks with a third octet of 255 within `MERIT-NT` and `SURF-NT`, and found that these blocks received 1.70–3.06% fewer sources relative to the median number of sources in /24 blocks in `MERIT-NT` across samples, and between 4.74% fewer and 0.45% more in `SURF-NT`. In contrast, this trend did not hold for `UCSD-NT`, whose –.–.255/24 blocks saw a slight increase in source counts, between 0.46% and 1.75% across samples.

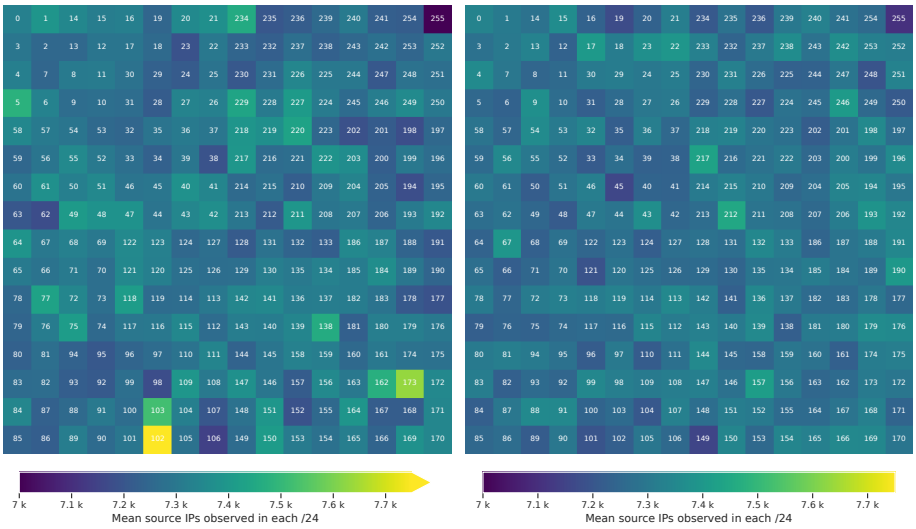

(a) MERIT-NT

(b) SURF-NT

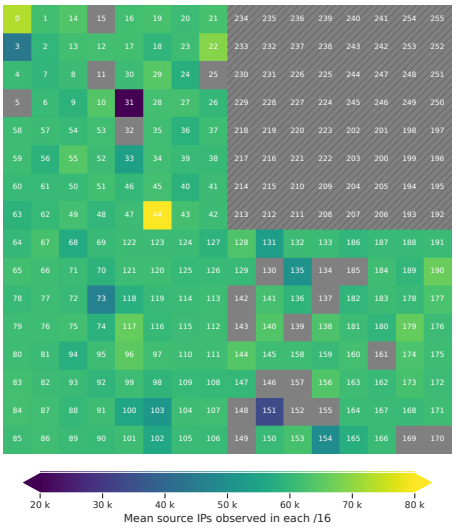

(c) UCSD-NT

Fig. 6: Number of source IP addresses observed per destination IP address block. Each cell represents a /24 block for MERIT-NT and SURF-NT, and a /16 block for UCSD-NT, labeled with the third and second octet of the block, respectively. Gray blocks indicate that no packets were observed (i.e., the subnet is assigned and therefore excluded), while the hatched region indicates the /10 prefix no longer part of UCSD-NT. An arrow at either end of a color bar indicates that values extend beyond the scale. Most blocks in each plot observed similar numbers of unique source addresses.

A few other blocks stood out. In MERIT-NT, the –.–.102/24 block within the first /16 block received 5.32–13.3% more sources than the median across all /24 blocks in that /16, although other /24 hotspots appeared in different /16s. In UCSD-NT, the –.44/16 block observed 32.8–92.2% more sources than the median across samples, whereas the –.31/16 block saw a substantial decrease, with 91.4–93.2% fewer sources than the median. The decreased sources observed in –.31/16 can be attributed to 99.2% of the address block being assigned and excluded from our dataset throughout the observation period. We no longer observed elevated source counts caused by BitTorrent and Conficker in the first /9 of UCSD-NT described in the original study [4].

IP Address Structure IBR sources may target addresses ending in .0 or .255 differently, as they may assume these addresses to be reserved network and broadcast addresses [23]. This is an important consideration when deploying telescopes on small address ranges. We counted the unique source addresses observed per last octet across all fully covered /24 telescope blocks.

We found that telescope addresses ending in .255 were generally targeted less than addresses with another final octet. In MERIT-NT, such addresses received between 0.18% more and 5.45% fewer sources than the median across samples. In SURF-NT, they were 4.30–7.71% less likely to be targeted, while in UCSD-NT they were 4.17–10.87% less likely. Earlier work [20] reported a similar pattern in MERIT-NT, although with large variations by port.

Addresses ending in .0 showed a less consistent pattern. In MERIT-NT and SURF-NT, such addresses observed 3.53–4.22% and 1.76–3.37% fewer sources addresses than the median across samples, respectively. In UCSD-NT, the trend was opposite with addresses with a last octet of 0 being targeted by 14.8–59.9% *more* sources. The previously identified anomalous –.44/16 block was the largest contributor to this effect. Even after excluding this block, UCSD-NT destination addresses ending in .0 remained 2.94–30.1% more likely to be targeted.

This discrepancy may stem from the mixed usage of the UCSD-NT space. The three UCSD-NT/16 blocks contributing most source addresses all contained assigned subnets in the 12 months prior to the observation period, suggesting potential local scanning activity or a memory effect.

6.6 Traffic Locality

As established in Section 6.1, telescope blocks of the same size generally observed similar numbers of source addresses. We investigated the extent to which IBR is specific to each telescope. The authors of the original study [4] compared traffic to MERIT-NT and UCSD-NT addresses with overlapping second, third, and fourth octets. However, because the three telescopes in our study do not share a common second octet (nor first octet), we randomly selected /16 blocks with different second octets[6] for the comparison to avoid inflated overlap caused by

[6] 83, 65, and 95 for UCSD-NT, MERIT-NT, and SURF-NT, respectively.

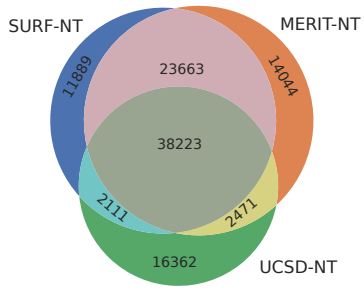

Fig. 7: Overlap in source IP addresses observed in /16 blocks across the three telescopes on April 5, 2025 at 3:00–4:00 UTC. The overlap between `SURF-NT` and `MERIT-NT` was consistently higher than that between either of them and `UCSD-NT`, while the overlap among all three telescopes was the consistently highest across samples.

IBR bias toward specific octets. We ensured that the `UCSD-NT`/16 block was not assigned and thus completely covered by the telescope. Figure 7 shows the overlap in source addresses for the first one-hour sample of the three telescopes. Across the samples we tested, the proportion of addresses observed in all three telescopes ranged from 35.1% to 39.4%.

In addition to the selected /16 blocks, we computed a Jaccard similarity matrix for all possible /16 combinations. We observed that the overlap between `MERIT-NT` and `SURF-NT` was consistently higher than that between `UCSD-NT` and either `MERIT-NT` or `SURF-NT`. This suggests that some IBR sources tend to not target `UCSD-NT` or specifically target both `MERIT-NT` and `SURF-NT`. Furthermore, the overall high overlap across all telescopes suggests that individual IBR sources were ubiquitous throughout the address space, indiscriminately transmitting at a sufficiently high rate to be observed in one-hour samples across three distributed /16s. Specifically, assuming uniformly random probing of the IPv4 address space, a Poisson approximation shows that a probing rate of at least 74.2 pps is required to reach all three /16s with 95% probability.

6.7 IP vs. Port Concentration

In addition to the analyses presented in the original study [4], we assess the number of destination addresses and ports sources targeted by IBR sources. We counted the destinations observed in the three /16 blocks of the three telescopes, using the same second octets that we analyzed in §6.6. Figure 8 shows the distribution of the number of destination addresses and port/protocol combinations contacted per source address. For the `MERIT-NT` and `SURF-NT` /16s, approximately 30.1% and 30.3% of sources contacted only one address of the telescope, and 57.1% and 57.6% contacted only one port. In contrast, the IBR observed in the `UCSD-NT` /16 showed a reversed pattern: 58.4% of sources contacted only one address, while just 36.8% contacted a single port. We found that this pattern

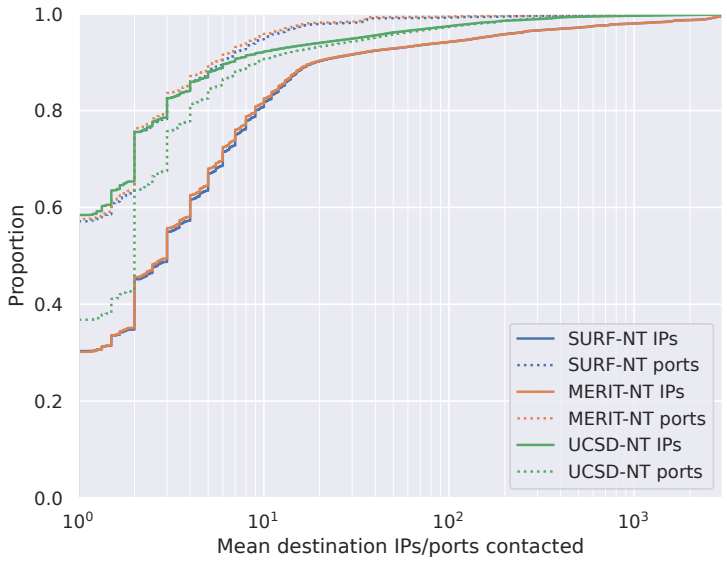

Fig. 8: ECDF of the mean destination addresses and protocol/ports combinations contacted per source address within /16 blocks across one-hour samples. Sources to `SURF-NT` and `MERIT-NT` contact more addresses than ports, whereas the opposite holds for `UCSD-NT`.

was consistent across all fully covered /16s of each telescope. A possible explanation for this difference is `UCSD-NT`'s much larger address space, which may cause scanners to adapt their probing strategies. Understanding this difference in targeting behavior warrants additional study.

6.8 RSDoS Detection

As a practical illustration of how the size of telescope impacts its utility, we considered its ability to infer RSDoS attacks. To this end, we follow the methodology of Moore et al. [24] (§2.2) on increasingly smaller portions of the `UCSD-NT` address space. We set the same parameters to the values that CAIDA uses to generate the RSDoS attack metadata feed [8]: at least 30 pps calculated over a 60-second sliding window, targeting 25 or more telescope addresses. We started by calculating the number of inferred RSDoS targets from the entire `UCSD-NT` address space (/9+/10), then removed the /10 block, and subsequently step-wise halved the telescope block size until we arrived at /16 blocks.

Figure 9 shows the median RSDoS targets inferred from one-hour samples per address block size. As the address space shrinks, the number of targets declines logarithmically. However, there is substantial variation in the number of targets inferred at the block level. For each block size and sample, we calculated the coefficient of variation (CV) of the inferred target counts across blocks. The

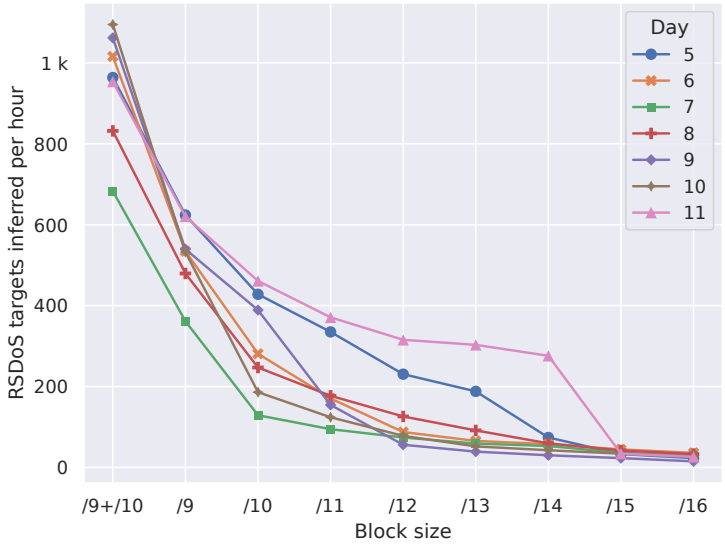

Fig. 9: Median RSDoS targets inferred from one-hour `UCSD-NT` samples. The observation counts decrease for larger IP address blocks. RSDoS target visibility decreased logarithmically with the size of the address space.

CVs ranged from 0.26 to 1.65, suggesting that spoofed source addresses are not uniformly random. Further investigation is required to confirm this observation.

Next, we compared the largest consecutive prefixes in each telescope, `UCSD-NT` (/9), `MERIT-NT` (/14), and `SURF-NT` (/16). A mean of 95.4% of RSDoS targets observed in the `MERIT-NT` block during one-hour samples were also detected in the `UCSD-NT` block during the same period, which aligns with findings from prior research [18]. Conversely, a mean of 13.2% of targets seen in `UCSD-NT` were also seen in `MERIT-NT`, although the latter comprised only 3.28%[7] of the address space of the former. Similarly, 93.9% of targets observed in `SURF-NT` were also detected in `UCSD-NT`, while 13.1% of targets seen in `SURF-NT` were also seen in `UCSD-NT`. However, in terms of the number of addresses, the `SURF-NT` block comprises only 0.8% of the `UCSD-NT` block. The relatively high overlap can be attributed to the randomly spoofed nature of such attacks and the fact that they rely on high packet rates to be effective. In general, using a distributed telescope makes it more difficult for attackers to evade all its address blocks.

Figure 10 compares the attacks that were observed in both `UCSD-NT` and `MERIT-NT`. The diagonal line indicates the proportional difference in monitored address space. We found that the number of packets inferred to be part of RSDoS attacks closely follows this line. Additionally, we found that some targets sent over five orders of magnitude more backscatter than others. The number of

[7] Only considering addresses within the /9 block that are monitored by `UCSD-NT`.

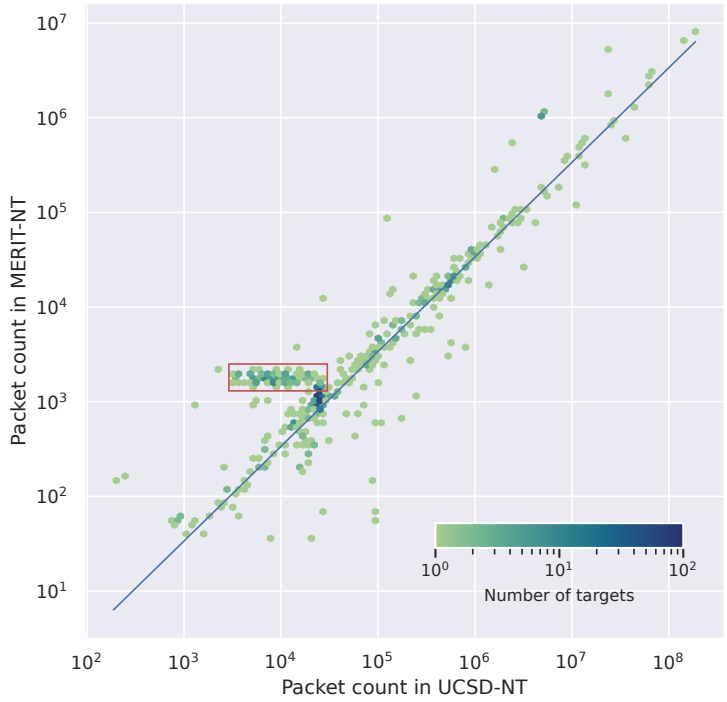

Fig. 10: Targets observed in `MERIT-NT` and `UCSD-NT` with corresponding backscatter packet counts. The color gradient indicates the number of targets observed for each packet count. The diagonal line represents the expected packet count based on the relative sizes of the telescopes. The majority of points within the rectangle belong to a clustered attack on a /24 subnet during a one-hour sample.

packets enables inference of a lower bound — packets may be dropped at or on the path to or from the target — on the packet rate of the attack.

One notable outlier is the cluster of targets demarcated by the rectangle. Upon inspection, 105 (76.1%) of these targets were inferred on April 5, 3:00–4:00 UTC. All of these 105 target addresses were targeted within this hour, they all fell within the same /24 subnet of an Argentinian game hosting provider. Furthermore, the drift of this cluster to the left of the diagonal indicates a skew in spoofed source addresses, which could be the result of a random number generator biased toward `MERIT-NT` addresses.

7 Discussion

Growing IPv4 shortage and rising operational costs prompt operators of network telescopes to reconsider their darknet deployments. Prior work [16,29,5,20,26,35] has proposed alternative approaches. These developments raise a central question: are large network telescopes still necessary as a data source for Internet-wide measurements?

Smaller, fragmented telescopes can provide viable alternatives. When measured over longer periods, smaller telescopes still achieve high visibility, i.e., the number of sources observed (§6.1), and their marginal visibility tends to be higher (§6.4, §6.6). Moreover, the generally even distribution of sources over blocks within large and smaller telescopes (§6.5) suggests that even limited portions of address space can capture representative IBR.

Nevertheless, large telescopes continue to provide a unique perspective. Our analysis showed that IBR sources exhibited different targeting patterns toward the larger telescope in our study (§6.2, §6.7). Furthermore, having a larger monitored address space enhances observation of ASes originating fewer addresses (§6.3) and lower-rate RSDoS attacks (§6.8). These results indicate that large-scale telescopes can observe events that smaller telescopes would likely miss.

The trade-off between operating large and smaller telescopes leaves operators with a practical dilemma: to downsize or not. Apart from the availability of address space, dimensioning decisions should consider the computation and storage requirements. Given the pressure to use IPv4 allocations efficiently, we recommend aligning telescope configurations with the study objectives. For example, research investigating topological scanning behavior would benefit from a telescope composed of many fragmented address blocks, whereas a study focusing on RSDoS attacks is better served by a telescope encompassing a large address space, regardless of its contiguity. While sub-/24 telescopes may suffice for certain studies, our observations show that addresses differing in their last octet are targeted differently, which can introduce bias if the selection is too narrow.

A collection of fragmented address blocks thus provide a cost-effective alternative to large, contiguous telescope ranges due to their high marginal visibility and potentially topological diversity. Analogous to radio astronomy, where dis-

tributed arrays provide a broader field of view, distributed network telescopes can monitor diverse segments of the global address space. Nevertheless, the increased management and synchronization complexity of such distributed deployments should also be taken into account.

7.1 Future Work

Filtering Spoofed Traffic While we made a best-effort attempt to filter bursty spoofing activity, we acknowledge that lower-rate spoofing activity may remain. Failing to filter such traffic could lead researchers to overestimate the number of sources observed and derived metrics, such as /24 blocks and ASNs. Another concern is that short-term spoofing activity effectively constitutes a denial-of-service attack on the telescope infrastructure and may create gaps in the captured packet traces. Future work should explore automatic derivation of heuristics to reliably and durably identify spoofed packets directly targeting a telescope.

Ephemeral Telescopes Leased address blocks and on-demand cloud-based sensors could increase geographical and topological coverage, with the additional benefit of being more difficult for malicious actors to avoid. However, their data may be biased by traffic intended for a previous tenant of the address space and thus requires novel filtering strategies [26].

8 Conclusions

The work by Benson et al. [4] examined how unsolicited traffic (called IBR) can be used to study a variety of Internet phenomena. We revisited their seminal work 10 years later, amid IPv4 address scarcity and surging IBR levels, and found overall consistent results. Our findings show that IBR continues to be a useful data source of Internet-wide measurements, despite being collected from telescopes covering a substantially smaller address space.

We also showed that although visibility increases with telescope size, it exhibits diminishing returns. This pattern held consistently throughout our study: the IBR sources, their originating ASes, and RSDoS targets all grew sublinearly with the size of the monitored space. Furthermore, we observed ubiquitous presence of IBR sources across telescopes, suggesting pervasive Internet-wide probing. Nevertheless, the largest telescope in our study, `UCSD-NT`, observed more sources in absolute terms and detected RSDoS events missed by the other two telescopes.

Operators can apply our methodology to guide telescope dimensioning decisions. Our findings suggest that distributed telescope deployments may provide a cost-effective alternative for maintaining broad visibility. Further research is needed to better understand spoofed traffic and environment-specific targeting behavior, as both factors can substantially influence what a telescope observes.

Acknowledgements. We thank our shepherd and the anonymous reviewers for their insightful comments. We also extend our gratitude to Merit and SURF for providing the telescope data used in this study.

This material is based on research sponsored by the National Science Foundation (NSF) grants CNS-2120399, CNS-2450552, OAC-2319959, OAC-2531134, by the European Commission under the GN5-2 project, by the Dutch Research Council (NWO) CATRIN project (NWA.1215.18.003), and by the Netherlands Enterprise Agency (RVO) MISD project under the 8ra initiative (IPCEI-CIS). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the funding agencies.

Ethical Considerations

This work does not raise ethical concerns.

References

1. Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., Cochran, J., Durumeric, Z., Halderman, J.A., Invernizzi, L., Kallitsis, M., Kumar, D., Lever, C., Ma, Z., Mason, J., Menscher, D., Seaman, C., Sullivan, N., Thomas, K., Zhou, Y.: Understanding the Mirai Botnet. In: 26th USENIX Security Symposium (USENIX Security 17). pp. 1093–1110 (2017), https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/antonakakis
2. Bailey, M., Cooke, E., Jahanian, F., Nazario, J., Watson, D.: The Internet Motion Sensor: A Distributed Blackhole Monitoring System. In: NDSS Symposium 2005 (Feb 2005), https://www.ndss-symposium.org/ndss2005/internet-motion-sensor-distributed-blackhole-monitoring-system/
3. Benson, K., Dainotti, A., Claffy, KC., Aben, E.: Gaining insight into AS-level outages through analysis of Internet Background Radiation. In: 2013 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS). pp. 447–452 (Apr 2013). https://doi.org/10.1109/INFCOMW.2013.6562915
4. Benson, K., Dainotti, A., claffy, k., Snoeren, A.C., Kallitsis, M.: Leveraging Internet Background Radiation for Opportunistic Network Analysis. In: Proceedings of the 2015 Internet Measurement Conference. pp. 423–436. IMC '15, Association for Computing Machinery, New York, NY, USA (Oct 2015). https://doi.org/10.1145/2815675.2815702
5. Bortoluzzi, F., Irwin, B., Beiler, L.S., Westphall, C.M.: Cloud Telescope: A distributed architecture for capturing Internet Background Radiation. In: 2023 IEEE 12th International Conference on Cloud Networking (CloudNet). pp. 77–85 (Nov 2023). https://doi.org/10.1109/CloudNet59005.2023.10490018
6. Camargo, A.V.C., Granville, L., Bertholdo, L.M.: Beyond Size: Investigating the Impact of Scaled-Down Network Telescopes on Threat Detection. International Journal of Network Management **35**(3), e70014 (2025). https://doi.org/10.1002/nem.70014
7. Center for Applied Internet Data Analysis (CAIDA): FlowTuple - STARDUST (Mar 2021), https://www.caida.org/projects/stardust/docs/data/flowtuple/

8. Center for Applied Internet Data Analysis (CAIDA): Randomly and Uniformly Spoofed Denial-of-Service (RSDoS) Attack Metadata (2025), https://www.caida.org/catalog/datasets/rsdos-targets/
9. Chindipha, S.D., Irwin, B., Herbert, A.: Quantifying the Accuracy of Small Subnet-Equivalent Sampling of IPv4 Internet Background Radiation Datasets. In: Proceedings of the South African Institute of Computer Scientists and Information Technologists 2019. pp. 1–8. SAICSIT '19, Association for Computing Machinery, New York, NY, USA (Sep 2019). https://doi.org/10.1145/3351108.3351129
10. corsaro3 contributors: CAIDA/corsaro3. CAIDA (2025), https://github.com/CAIDA/corsaro3
11. Dainotti, A., Amman, R., Aben, E., Claffy, K.C.: Extracting benefit from harm: Using malware pollution to analyze the impact of political and geophysical events on the internet. ACM SIGCOMM Computer Communication Review **42**(1), 31–39 (Jan 2012). https://doi.org/10.1145/2096149.2096154
12. Dainotti, A., Benson, K., King, A., claffy, k., Kallitsis, M., Glatz, E., Dimitropoulos, X.: Estimating internet address space usage through passive measurements. ACM SIGCOMM Computer Communication Review **44**(1), 42–49 (Dec 2014). https://doi.org/10.1145/2567561.2567568
13. Dainotti, A., King, A., Claffy, k., Papale, F., Pescapè, A.: Analysis of a "/0" stealth scan from a botnet. In: Proceedings of the 2012 Internet Measurement Conference. pp. 1–14. IMC '12, Association for Computing Machinery, New York, NY, USA (Nov 2012). https://doi.org/10.1145/2398776.2398778
14. Durumeric, Z., Bailey, M., Halderman, J.A.: An Internet-Wide View of Internet-Wide Scanning. In: 23rd USENIX Security Symposium (USENIX Security 14). pp. 65–78 (2014), https://www.usenix.org/node/184494
15. Gao, M., Mok, R., Carisimo, E., Li, E., Kulkarni, S., claffy, k.: DarkSim: A similarity-based time-series analytic framework for darknet traffic. In: Proceedings of the 2024 ACM on Internet Measurement Conference. pp. 241–258. IMC '24, Association for Computing Machinery, New York, NY, USA (Nov 2024). https://doi.org/10.1145/3646547.3688426
16. Harrop, W., Armitage, G.: Defining and Evaluating Greynets (Sparse Darknets). In: The IEEE Conference on Local Computer Networks 30th Anniversary (LCN'05)l. pp. 344–350 (Nov 2005). https://doi.org/10.1109/LCN.2005.46
17. Hiesgen, R.: Spoki: Unveiling a New Wave of Scanners through a Reactive Network Telescope. In: 31st USENIX Security Symposium (USENIX Security 22) (2022), https://www.usenix.org/conference/usenixsecurity22/presentation/hiesgen
18. Hiesgen, R., Nawrocki, M., Barcellos, M., Kopp, D., Hohlfeld, O., Chan, E., Dobbins, R., Doerr, C., Rossow, C., Thomas, D.R., Jonker, M., Mok, R., Luo, X., Kristoff, J., Schmidt, T.C., Wählisch, M., claffy, k.: The Age of DDoScovery: An Empirical Comparison of Industry and Academic DDoS Assessments. In: Proceedings of the 2024 ACM on Internet Measurement Conference. pp. 259–279. IMC '24, Association for Computing Machinery, New York, NY, USA (Nov 2024). https://doi.org/10.1145/3646547.3688451
19. Internet Assigned Numbers Authority (IANA): Service Name and Transport Protocol Port Number Registry (May 2025), https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml
20. Izhikevich, L., Tran, M., Kallitsis, M., Fass, A., Durumeric, Z.: Cloud Watching: Understanding Attacks Against Cloud-Hosted Services. In: Proceedings of the 2023 ACM on Internet Measurement Conference. pp. 313–327. IMC '23, Association for

Computing Machinery, New York, NY, USA (Oct 2023). https://doi.org/10.1145/3618257.3624818
21. Jonker, M., King, A., Krupp, J., Rossow, C., Sperotto, A., Dainotti, A.: Millions of targets under attack: A macroscopic characterization of the DoS ecosystem. In: Proceedings of the 2017 Internet Measurement Conference. pp. 100–113. IMC '17, Association for Computing Machinery, New York, NY, USA (Nov 2017). https://doi.org/10.1145/3131365.3131383
22. Merit: Orion Network Telescope (2025), https://www.merit.edu/research/projects/orion-network-telescope/
23. Mogul, J., Postel, J.: Internet standard subnetting procedure. RFC 950, IETF (Aug 1985), http://tools.ietf.org/rfc/rfc950.txt
24. Moore, D., Shannon, C., Brown, D.J., Voelker, G.M., Savage, S.: Inferring Internet denial-of-service activity. ACM Transactions on Computer Systems **24**(2), 115–139 (May 2006). https://doi.org/10.1145/1132026.1132027
25. Pang, R., Yegneswaran, V., Barford, P., Paxson, V., Peterson, L.: Characteristics of internet background radiation. In: Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement. pp. 27–40. IMC '04, Association for Computing Machinery, New York, NY, USA (Oct 2004). https://doi.org/10.1145/1028788.1028794
26. Pauley, E., Barford, P., McDaniel, P.: DScope: A Cloud-Native Internet Telescope. In: 32nd USENIX Security Symposium (USENIX Security 23). pp. 5989–6006 (2023), https://www.usenix.org/conference/usenixsecurity23/presentation/pauley
27. Prehn, L., Lichtblau, F., Feldmann, A.: When wells run dry: The 2020 IPv4 address market. In: Proceedings of the 16th International Conference on Emerging Networking EXperiments and Technologies. pp. 46–54. CoNEXT '20, Association for Computing Machinery, New York, NY, USA (Nov 2020). https://doi.org/10.1145/3386367.3431301
28. Richter, P., Allman, M., Bush, R., Paxson, V.: A Primer on IPv4 Scarcity. ACM SIGCOMM Computer Communication Review **45**(2), 21–31 (Apr 2015). https://doi.org/10.1145/2766330.2766335
29. Richter, P., Berger, A.: Scanning the Scanners: Sensing the Internet from a Massively Distributed Network Telescope. In: Proceedings of the Internet Measurement Conference. pp. 144–157. IMC '19, Association for Computing Machinery, New York, NY, USA (Oct 2019). https://doi.org/10.1145/3355369.3355595
30. RouteViews: University of Oregon RouteViews Project (May 2025), https://www.routeviews.org/routeviews/
31. Sommese, R., Claffy, KC., van Rijswijk-Deij, R., Chattopadhyay, A., Dainotti, A., Sperotto, A., Jonker, M.: Investigating the impact of DDoS attacks on DNS infrastructure. In: Proceedings of the 22nd ACM Internet Measurement Conference. pp. 51–64. IMC '22, Association for Computing Machinery, New York, NY, USA (Oct 2022). https://doi.org/10.1145/3517745.3561458
32. Soro, F., Drago, I., Trevisan, M., Mellia, M., Ceron, J., J. Santanna, J.: Are Darknets All The Same? On Darknet Visibility for Security Monitoring. In: 2019 IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN). pp. 1–6 (Jul 2019). https://doi.org/10.1109/LANMAN.2019.8847113
33. Tanveer, H.B., Chan, W.S., Mok, R.K.P., Kappes, S., Richter, P., Gasser, O., Ronan, J., Berger, A., Claffy, k.: Unveiling IPv6 Scanning Dynamics: A Longitudinal Study Using Large Scale Proactive and Passive IPv6 Telescopes (Aug 2025). https://doi.org/10.48550/arXiv.2508.07506

34. The Number Resource Organization: RIR Statistics (Apr 2025), https://www.nro.net/about/rirs/statistics/
35. Wagner, D., Ranadive, S.A., Griffioen, H., Kallitsis, M., Dainotti, A., Smaragdakis, G., Feldmann, A.: How to Operate a Meta-Telescope in your Spare Time. In: Proceedings of the 2023 ACM on Internet Measurement Conference. pp. 328–343. IMC '23, Association for Computing Machinery, New York, NY, USA (Oct 2023). https://doi.org/10.1145/3618257.3624831
36. Wustrow, E., Karir, M., Bailey, M., Jahanian, F., Huston, G.: Internet background radiation revisited. In: Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement. pp. 62–74. IMC '10, Association for Computing Machinery, New York, NY, USA (Nov 2010). https://doi.org/10.1145/1879141.1879149