

# Towards validation of the Internet Census 2012

Dirk Maan, José Jair Santanna, Anna Sperotto, and Pieter-Tjerk de Boer

Design and Analysis of Communication Systems (DACS)  
University of Twente, The Netherlands  
h.c.maan@student.utwente.nl  
{j.j.santanna,a.sperotto,p.t.deboer}@utwente.nl

**Abstract.** The reliability of the “Internet Census 2012” (IC), an anonymously published scan of the entire IPv4 address space, is not a priori clear. As a step towards validation of this dataset, we compare it to logged reference data on a /16 network, and present an approach to systematically handle uncertainties in timestamps in the IC and reference data. We find evidence the scan indeed took place, and a 93% match with the /16 reference data.

**Keywords:** Internet census, scan, validation.

## 1 Introduction

In March 2013, an anonymous researcher published the result of a project called *Internet Census 2012* (IC) [2]. The project was based on a scan of the entire IPv4 addresses space (i.e., an Internet-wide scan). A scan is created by sending probe packets to hosts, using one probing host or a distributed network of hosts controlled by a central server. Technical issues that could prevent accurate results are packet loss, or bot misconfiguration. The anonymous author claims his 9 Tbytes of raw log files are the most recent and accurate census of the Internet.

Active scans on the Internet are certainly an important source of information, as several studies have demonstrated that they can help reveal new kinds of vulnerabilities, monitor deployment of mitigation, and highlight hidden distributed ecosystems [5, 6, 7, 8, 11]. However, the IC results were published anonymously, and the methodology only partially described, which, as pointed out by the Cooperative Association for Internet Data Analysis (CAIDA) [1], leads to some important questions, such as: how does one know that the IC scan actually happened, and if it did, how does one know that the resulting data is correct?

In this paper we extend the work of CAIDA [1] by proposing a methodology to validate the Internet Census in third-party networks and datasets. We then applied our methodology to the /16 network address block of the University of Twente (UT), and found a match of about 93% between the considered Internet Census data and our reference data.

The scientific community has ethical (and legal) concerns about the IC. There are concerns about network-wide scanning in general; discussions so far have not led to clear consensus. In the case of the IC, this is exacerbated by the fact that

## 2. RELATED WORK AND BACKGROUND

the IC was performed using a botnet consisting of around 420 thousand compromised systems. In this paper, we sidestep these concerns by limiting ourselves to comparing this dataset to reference data of our own, with the sole purpose of finding out to what extent the IC dataset actually reflects reality.

The remainder of this paper is organized as follows. Section 2 discusses related work, and is followed by a characterization of the IC dataset in Section 3. Section 4 describes our proposed methodology, which is analyzed in Section 5. Finally, in Section 6 we summarize our findings and highlight future work.

## 2 Related work and Background

Since the beginning of the Internet, scans have been performed to obtain information about the end hosts. RFC-832 [12] describes the first documented scan of the Internet. At that time, in 1982, 315 hosts were probed to see if they use the TCP protocol. The scan took roughly one day.

In [6], Heidemann et al. presented a study of the active Internet over the period 2003–2008. The presented census highlighted anomalies in the un-allocated address space and indicates the percentage of usage for allocated network blocks. The observations from this scan were validated by comparing them to scans of smaller address blocks. This study only considers ICMP probes, as TCP is considered too resource-consuming for the scanned hosts.

Furthermore, Holz et al. conducted HTTPs scans of the top million popular hosts over a timespan of 1.5 years [9]. These scans were horizontal, as only port 443 was probed, scanning for certificates. The IC differs by probing the top 100 and several other random ports.

Another Internet-wide scan which is similar to the IC was performed by Durumeric et al. [4]. To the best of our knowledge, it is the most recent documented scan of the complete Internet. The authors developed a scanning tool, ZMap, specifically designed to perform fast scans at a large scale. Differently from the IC, the scans performed by Durumeric et al. are targeted at the study of specific protocols, showing that the authors are cautious to avoid more scanning activities than needed.

Internet-wide scans are unfortunately not only used for Internet measurements. In [3], Dainotti et al. describe a scan of the entire IPv4 address space performed by the Salinity botnet. It is estimated that the botnet has scanned approximately 3 million distinct IP addresses over a period of 12 days, scanning both port 5060 (SIP) and port 80.

## 3 Characterization of Internet Census 2012

The dataset provided by the IC is composed of seven sets of traces, as summarized in Table 1. Each entry of each trace contains three elements: 1) the IP address of the probed device, 2) a timestamp indicating the moment of probing, and 3) the result of the scan, which depends on the scan method used in each

### 3. CHARACTERIZATION OF INTERNET CENSUS 2012

**Table 1.** Traces in IC

#	Trace	Content	Size
1	ICMP Ping	Responsiveness and latencies	1.8 TB
2	Reverse DNS	DNS records	366 GB
3	Serviceprobes	Services behind open ports	5.5 TB
4	Hostprobes	Responsiveness	771 GB
5	Syncscan	State of ports	435 GB
6	TCP/IP Fingerprint	Type of device and operating system	50 GB
7	Traceroute	Path of data packet	18 GB

trace. For example, the ICMP scan indicates if a host is reachable, while the Syncscan trace lists the status of the scanned ports.

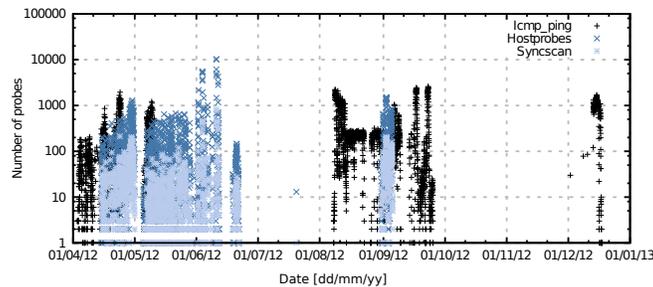
In our research, the traces ‘*icmp\_ping*’, ‘*hostprobes*’ and ‘*syncscan*’ are of special interest, because these traces indicate if a device was active or not at a certain timestamp. These three traces will be the only traces considered in the following sections.

#### 3.1 Trace overview

Figure 1 shows the time distribution of the probes that reached the UT /16 netblock accordingly to the IC. They are clustered in three main periods, namely April–July 2012, August–October 2012 and mid-December 2012. IP address are generally probed more than once per trace, as shown in Figure 2. The x-axis of this figure is the number of probes sent to an IP address, while the y-axis the respective frequencies. For example, in the hostprobes trace more than 35000 IP addresses were probed five times. In the syncscan, most IP addresses are probed only once or twice, while icmp\_ping has around 13 probes per IP address.

#### 3.2 Timestamp rounding

The IC paper does not much provide information about the probe timestamps reported in the datasets. Based on the format, the timestamps used in the IC



**Fig. 1.** Number of probes per day for 130.89.0.0/16

## 4. METHODOLOGY

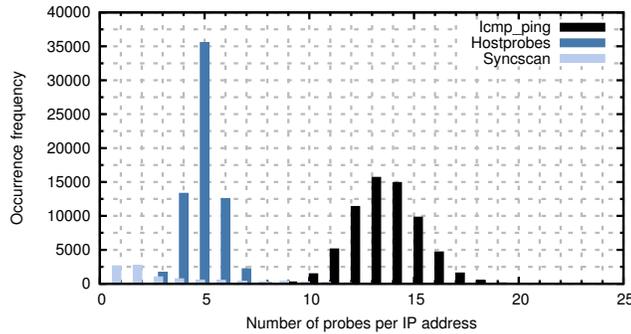


Fig. 2. Occurrence frequency of the number of probes per IP address

are assumed to be in standard Unix time format (i.e., seconds since Jan. 1, 1970). The source of the timestamping is unclear: they might be the probing bots themselves, or some central server collecting the data. Especially in the former case, timestamps may be off and inconsistent due to the respective bot’s clock not being set correctly.

Furthermore, it is notable that each timestamp value in the IC is an odd multiple of 900 seconds; in other words, each timestamp is either exactly 15 minutes before or after a full hour. Apparently, the actual timestamps of the probes have been rounded in some way, presumably for anonymization. The rounding strategy is not described; obvious possibilities are always up, always down, or to nearest. This leads to a total uncertainty of 3600 seconds, ranging from 1800 seconds before the IC timestamp to 1800 seconds after it. We will take this observation into account in the validation of the IC traces performed in Sections 4 and 5.

## 4 Methodology

Our validation methodology consists of two parts. First, we verify that probes from the IC have indeed reached hosts at the UT. We do this by analysing packet traces that were collected at the time of the scan for a particular IP address. Secondly, we validate the information of the IC for the /16 netblock of the UT, which we indicate as  $IC_{UT}$ , by comparing the IC traces with a reference dataset based on the ARP tables of the UT routers.

### 4.1 Single-IP analysis

For the single-IP analysis, the trace of all incoming traffic to a server at the UT is compared with the IC data. (We only had such a trace available for a single machine.) The best trace of the IC to use for this goal is the syncscan trace, because of its many probes in a short time interval. The hostprobes and icmp\_ping IC traces are not inspected, because, due to the timestamp rounding

**Table 2.** Overall comparison of the *ARP* and *IC<sub>UT</sub>* datasets

Subset	Definition	Description
Subset A	$ARP - IC_{UT}$	UT IP addresses are active but not included in the IC
Subset B	$ARP \cap IC_{UT}$	UT IP addresses are active and included in the IC
Subset C	$IC_{UT} - ARP$	UT IP addresses are not active but included in the IC

and the frequency of ICMP packets in the host trace, it was not possible to match ICMP packets in the trace to the IC with certainty. From the syncscan trace the probed ports and the time of probing are known. This is enough data to filter the IC probes from other incoming traffic at the server. To validate each entry that corresponds to the analysed IP, we consider the following: first, it is checked if the machine was actually approached by the IC at the stated timestamp; second, the state that was reported in the IC is checked to be equal to the actual state of the machine.

#### 4.2 UT address block analysis

**Reference dataset** To validate the IC data, we need a reference dataset indicating, for a certain moment in time, which IP addresses are active in the /16 UT networks. The dataset used in this paper is a consolidated snapshot of the ARP table of the UT routers, referred as *ARP*, during the period April 2012–December 2012. From the analysis of the ARP tables we determine that the UT /16 block has a utilization of about 53%. We are aware that the ARP tables can introduce some measurement imprecisions. For example, an IP would typically remain in the ARP tables for some time after it has disconnected. Furthermore, several gaps in the dataset are present, due to SNMP timeouts that occurred or because the database table space was temporarily full. However, considering that we are analyzing IC data over a period of 8 months, we believe that these imprecisions only have a limited impact on the validation.

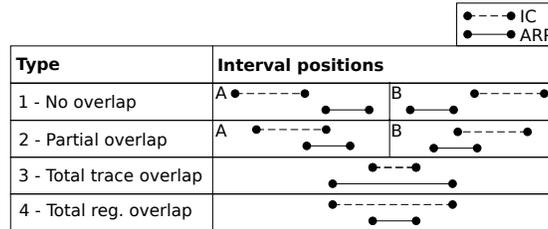
We first investigate the intersection and, respectively left and right difference between the IP sets in *IC<sub>UT</sub>* and *ARP*, as indicated in Table 2. Subsets A and C report errors in the IC, namely active IP addresses that have not been reported or, conversely, inactive IP addresses that have been wrongly included in the IC. In the case of Subset B, we proceed as described in the following subsection.

**Subset B timestamp analysis** Although in principle Subset B is the intersection of *IC<sub>UT</sub>* and *ARP*, this is not sufficient to state that these IP addresses are correctly reported. For example, an IP could have been listed in the IC at a moment in time in which it was not active, or viceversa. To investigate this issue, we perform the following analysis on the ARP and IC timestamps.

First, we consider the complete interval in which the IC probe can be sent. Due to the timestamp rounding in the IC, every timestamp  $t$  is expanded to create an interval with start time  $t - 1800$  sec and end time  $t + 1800$  sec. For the ARP tables this is not necessary, since they already report the start and end time

#### 4. METHODOLOGY

of activity for a certain IP. Comparing the interval of the ARP table with the interval of the IC will result in four possible outcomes, referred as overlaptypes (see Figure 3) with four different conclusions.



**Fig. 3.** Different overlap possibilities.

*No overlap:* This situation occurs when an IP address was registered according to the ARP table, but the IC probed the IP address outside the ARP interval. Therefore, the IC should not state this IP address as being alive.

*Partial overlap:* If the intervals overlap only partially, it is not clear whether the probe was sent within or outside the ARP registration interval, so no conclusion can be drawn.

*Total IC overlap:* In this case, the IC interval is completely contained in the ARP registration interval. The information in the IC must therefore contain an indication of the IP address being alive. If the IC states that the host is unreachable, this is an error.

*Total ARP overlap:* The last case is characterized by the ARP registration interval being completely overlapped by the IC trace interval. The IC trace interval itself is partially overlapped, so it is not clear whether the probe was sent within or outside the ARP registration interval and no conclusion can be drawn.

**Validation** In this study, we focus in particular on the erroneous entries in the IC, because they are an indication of the reliability of the IC. An erroneous entry is defined as an entry that contradicts the data in the ARP tables of the UT.

Basically, the comparison consists of the following steps, which will be elaborated further on:

- Split the IC trace into unreachable and alive subtraces; this step is necessary to correctly analyze the IPs in Subset B.
- Determine the appropriate subset for each entry.
- Count probes per subset.

In order to clarify the comparison in a more visual way, we refer the reader to the flowchart in Figure 4. This figure shows how an IC probe is categorized in a certain subset. Furthermore, it shows what type of IC probes are erroneous, as can be deduced from Section 4.2.

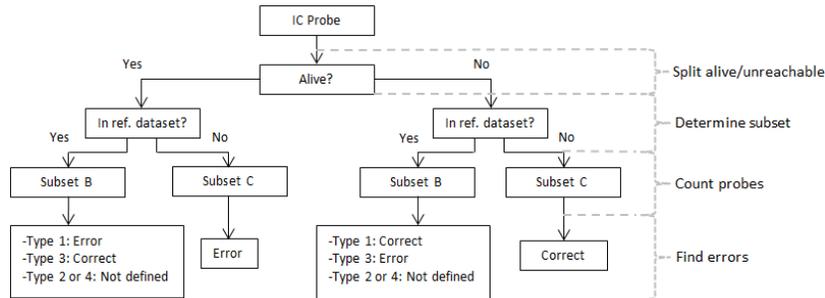


Fig. 4. Flowchart of the comparison process

## 5 Results

Similar to Section 4, the results are described in two parts and ordered in the same way, first the single-IP analysis is discussed, followed by the results of the /16 address block analysis.

### 5.1 Single-IP analysis results

Performing the method described in Section 4.1 on the given traces, results in the conclusion that the IC records concerning the server correspond to the traffic traces of our server. Indeed, the IP of the server turns out to be present in the IC records of the syncscan trace. Furthermore, in the incoming traffic of the server packets have been identified that match the timestamp and portnumbers listed for scans of this machine in the IC. As an aside, it turns out that when the IC scanned several ports within a short timeframe (same rounded timestamp), these scans came from the same IP address, while scans farther apart came from different IP addresses. Most probes were replied to with a packet having the RST and ACK flags set, revealing that the probed host exists but has these ports closed. These ports of the server IP were marked as ‘closed’ in the IC, which is as expected. One port was correctly marked as open in the IC. In the remaining cases, the IC marks ports for our IP as ‘filtered’, and indeed no corresponding incoming sync packet could be found.

By comparing the timestamps of the packet traces of the single host to the probe timestamps in the IC, the rounding of IC timestamps can be studied. It turns out that the timestamps in the IC are between 200 and 900 seconds lower than the timestamps in the packet trace. From this we can conclude that out of the three options mentioned in 3.2, only the round “always down” and “to nearest” are compatible with the data. (Note that we could distinguish between these remaining two options if we had timestamps in the first or third quarter of the hour, but apparently our server was only probed in the second and fourth quarter of the hour.) Although we only have this data of one IP and one trace, it is assumed that all timestamps in the IC are rounded the same way. In principle,

## 5. RESULTS

this reduces the timestamp uncertainty from 3600 seconds as discussed in 3.2, to 2700 seconds. However, in Section 4.2 we have used the original 3600 second uncertainty, both because 2700 s provides only a small advantage and for lack of time to redo the entire analysis.

By inspecting the packet traces of the host considered in this study, we observe the following:

- When the probed host replies by sending a packet with flags [RST,ACK] set, the IC reports a closed state accordingly.
- When the IC probe does not reach the probed host, the IC reports a filtered state accordingly.
- When the IC probe reaches the probed host and the host does not reply but drops the probe, the IC reports a filtered state accordingly.
- Some ports are probed multiple times.

Since each probe of our server in the IC, reporting a different state than filtered, indeed can be matched to a packet in our trace, we can conclude that the scan did indeed happen (confirming findings of [1]), and did indeed reach our network. Furthermore, we have obtained partial knowledge of the timestamp rounding.

### 5.2 /16 address block analysis results

In conformity with the division in subsets of Section 4.2, the results of the /16 address block analysis are split in three parts, because each of these parts requires a different analysis. The result of the interval comparison in subset B is an overview of the occurrence of different overlaptypes described in Figure 3.

**Subset A** All hosts registered in the ARP tables were present in `icmp_ping` and `hostprobes` traces, resulting in an empty subset A for these traces. By analysing the comparison of the `syncscan` trace, 71.4% of the IPs in `ARP` do not appear in the `ICUT`. Since the author of the IC paper states that a `syncscan` was limited to about 660 million IP addresses [2], we consider our observation in line with the IC description.

**Subset C** The percentage of the entries in subset C with respect to the total probes in each subtrace (e.g. `hostprobes_alive` or `icmp_unreachable`) is shown in the first column of Table 3.

According to the characteristics described in Section 4.2, no IC entries marked as alive should be in this subset. As shown in Table 3, the alive subtraces of `hostprobes`, `icmp_ping` and `syncscan` contain 5.31%, 2.82%, and 3.94% respectively nonmatching IP addresses. These are considered erroneous. However, further analysis shows that several alive IP addresses of trace `icmp_ping` that are categorized in this subset are actually broadcast addresses or network addresses. The UT has about 200 of these addresses in the /16 address block. These addresses

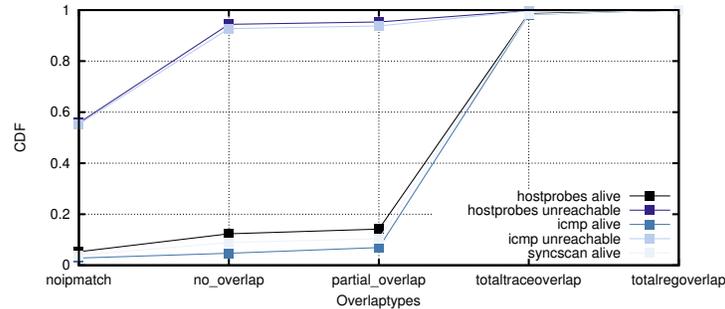
are categorized in subset C, because they are not included in the ARP table of the UT. The IC reported these addresses as alive, due to the reply of some UT routers that received the probe. In addition, several hosts reported as alive in this subset were probed close in time to moments where ARP table errors occurred. This can be another reason why some IP addresses in the /16 address block that were alive according to the IC are not present in the ARP table. From this observation is concluded that alive entries in this subset are not necessarily errors in the IC.

Figure 5 shows that more than 50% of the probes in the unreachable subtraces of hostprobes and icmp\_ping do not match the IP addresses of the ARP table. These probes are consistent with the ARP data, due to the fact that IP addresses can not be alive without them being registered in the ARP table. Hence these probes are considered correct.

**Subset B** By categorizing the records of subset B into separate overlaptypes, the validity of each subtrace can be determined. The result is a table with the number of probes that is counted for each overlaptype, as shown in Table 3.

**Table 3.** Result of comparison per subtrace of IC with ARP table

Subtrace	Subset C		Subset B				Total
	noipmatch		no_overlap	partial_overlap	totaltraceoverlap	totalregoverlap	
hostprobes_alive	1192 [5.31%]		1573 [7.01%]	411 [1.83%]	18945 [84.46%]	311 [1.39%]	22432
hostprobes_unreachable	170565 [55.80%]		118078 [38.63%]	2775 [0.91%]	13481 [4.41%]	769 [0.25%]	305668
icmp_alive	1497 [2.82%]		1010 [1.90%]	1176 [2.22%]	48406 [91.30%]	930 [1.75%]	53019
icmp_unreachable	458833 [55.27%]		310824 [37.44%]	8688 [1.05%]	49048 [5.91%]	2721 [0.33%]	830114
syncscan_alive	670 [3.94%]		839 [4.94%]	264 [1.55%]	14930 [87.90%]	282 [1.66%]	16985



**Fig. 5.** Cumulative Distribution Function of overlaptypes in subtraces

Figure 5 shows the same data as Table 3, but now as a cumulative distribution function. On the x-axis are the possible overlaptypes, ordered roughly by increasing quality of the match; the y-axis displays the percentage of probes having up to that kind of match. The following two paragraphs summarize the

## 6. CONCLUSIONS

observed distribution of the subtraces based on the probe states. Errors and correctness are discussed afterwards.

*Unreachable* As seen in Figure 5, the largest increase of the unreachable subtraces occurs in between overlaptypes `noipmatch` and `no_overlap`. Because these traces add up to more than 90% of all probes in these traces, we can conclude that most of these probes did indeed not reach active hosts. About these probed IP addresses it can be said that their unreachable state is noted correctly in the IC.

*Alive* Another observation from Figure 5 is the large percentage of the totaltraceoverlap in traces `hostprobes_alive`, `icmp_alive`, shown by a clear increase from overlaptypes `partial_overlap` to `totaltraceoverlap`. The share of total probes in these subtraces in overlaptypes `totaltraceoverlap` of subset B is 84.46% and 91.30% respectively. The probes of these traces in this category are correct. A similar percentage of 87.90% totaltraceoverlap is observed in the `syncscan_alive` subtrace.

**Trace correctness** In order to give an overall correctness indication for each IC trace for this /16 block, we take together the statistics for the unreachable and the alive subtraces of each trace. The cases which are counted as correct are marked green in Table 3: unreachable and either `noipmatch` or `no_overlap`, or alive and `totaltraceoverlap`. The cases which are counted as incorrect are marked red in the table: alive and either `noipmatch` or `no_overlap`, or unreachable and `totaltraceoverlap`. In the remaining cases, no correctness conclusion can be drawn due to the timestamp uncertainties. Thus, we cannot calculate the overall correctness percentage as a single value, but as a range to represent the uncertain cases.

A further difficulty occurs in the case of the `syncscan` trace: here we only have an alive subtrace and not an unreachable subtrace. This is because when a port was found not reachable, the IC reports it as ‘filtered’, indicating that the host may actually have been alive, but unreachable due to a filter on the path. Comparing this to our reference data of alive hosts clearly is not meaningful.

Finally, this results in the correctness ranges of each IC trace. Regarding `hostprobes`, the correctness is in between 93.75%–95.05%. For `icmp_ping` this correctness is equal to 92.63%–94.16% and in `syncscan`, 40.69%–95.89% is correct. The rather low lower bound in the latter case is because for we could only check the alive subtrace, as explained above.

## 6 Conclusions

In this paper, we have validated a /16 block from the “Internet Census 2012” by comparing it to locally logged data from that /16 block, and introduced a method to deal with the uncertainties in both the timestamps of the IC and the reference data.

## 6. CONCLUSIONS

Using the incoming traffic traces of a single host, it was validated that the IC scan included devices in the /16 address block of the UT. We were able to identify the synscan probes of the IC in the normal server traffic, which is a strong indication that the IC indeed was performed on the UT network. We showed that the synscan probe timestamps were either rounded down or to nearest with respect to the probe timestamp as observed on our server.

After verifying that the scan has taken place at the UT network, our analysis of the complete UT /16 address block has shown that about 93% of the IP addresses utilization in this block are correctly reported by the IC, by comparing them to logged ARP tables. Although this results indicates that the census has the possibility of depicting an accurate picture of the Internet utilization, it is also important to notice that, on a large scale this potentially amount to several millions of incorrectly classified hosts. Also, we reckon that the accuracy could be influenced by specific network settings, therefore we do not extrapolate from these results to wider conclusions.

Many error sources could have affected the IC scan when it was performed. Some of the possible error causes might be packet loss or bot misconfiguration. Missing information could for example lead to an incorrect unreachable state of an IP address in the IC. In contrast to the hostprobes and icmp\_ping traces, the synscan trace has only been validated for about 40%. Many probes of this trace were ignored in the process, e.g. UDP probes and probes that have the state ‘filtered’ in the IC. Therefore, we consider this result as not really accurate and think it should not be used as a measure for correctness of the entire synscan trace.

There are several opportunities for future work. By using traces of incoming traffic of more hosts, if available, the rounding method of the IC timestamps can be identified, and the accuracy of the timestamps studied further. If the rounding method is known, the accuracy of the proposed validation method can be increased. Furthermore, other address blocks of the IPv4 address space can be validated using the proposed method. When more address blocks in the IC are validated, a better conclusion about the validity of the entire IC can be drawn. Another possibility for future research is the validation of IC traces that were skipped in our research. The serviceprobes and tcp\_ip\_fingerprint traces for example contain Nmap data about the devices that were scanned.

**Acknowledgements** Special thanks goes to Jeroen van Ingen at ICTS (University of Twente) for providing the ARP data. This work was funded by the FP7 Network of Excellence project FLAMINGO (ICT-318488).

## References

1. Alistair, Alberto: Carna botnet scans confirmed. [http://blog.caida.org/best\\_available\\_data/2013/05/13/carna-botnet-scans/](http://blog.caida.org/best_available_data/2013/05/13/carna-botnet-scans/) (2013)
2. Anonymous: Internet census 2012. <http://internetcensus2012.bitbucket.org/paper.html> (2013)

## 6. CONCLUSIONS

3. Dainotti, A., King, A., Claffy, k., Papale, F., Pescapè, A.: Analysis of a "/" stealth scan from a botnet. In: Proceedings of the 2012 ACM Internet Measurement Conference. pp. 1–14. IMC '12, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2398776.2398778>
4. Durumeric, Z., Wustrow, E., Halderman, J.A.: Zmap: Fast internet-wide scanning and its security applications. In: 22nd USENIX Security Symposium. pp. 605–619 (2013)
5. Eckersley, P., Burns, J.: An observatory for the SSLiverse. Talk at Defcon 18 (2010), <https://www.eff.org/files/DefconSSLiverse.pdf>
6. Heidemann, J., Pradkin, Y., Govindan, R., Papadopoulos, C., Bartlett, G., Bannister, J.: Census and survey of the visible internet. In: Proceedings of the 2008 ACM Internet Measurement Conference. pp. 169–182. ACM, Vouliagmeni, Greece (2008), <http://www.isi.edu/~johnh/PAPERS/Heidemann08c.html>
7. Heninger, N., Durumeric, Z., Wustrow, E., Halderman, J.: Mining your ps and qs: Detection of widespread weak keys in network devices. In: 21st USENIX Security Symposium (2012)
8. Holz, R., Braun, L., Kammenhuber, N., Carle, G.: The SSL landscape: A thorough analysis of the X.509 PKI using active and passive measurements. In: 11th ACM SIGCOMM conference on Internet measurement (IMC) (2011)
9. Holz, R., Braun, L., Kammenhuber, N., Carle, G.: The SSL landscape: A thorough analysis of the x.509 PKI using active and passive measurements. In: Proceedings of the 2011 ACM Internet Measurement Conference. pp. 427–444. IMC '11, ACM, New York, NY, USA (2011), <http://doi.acm.org/10.1145/2068816.2068856>
10. Maan, H.C.: Validation of Internet Census 2012. BSc Thesis, University of Twente, The Netherlands (2014)
11. Moore, H.: Security flaws in universal plug and play. unplug. don't play (2013), <http://community.rapid7.com/servlet/JiveServlet/download/2150-1-16596/SecurityFlawsUPnP.pdf>
12. Smallberg, D.: RFC 832: Who talks TCP? (1982)