

# Classifying Internet One-way Traffic

TIK Report Number 336. May 2012.

Eduard Glatz  
Xenofontas Dimitropoulos  
Computer Engineering and Networks Laboratory  
ETH Zurich  
{eglatz,fontas}@tik.ee.ethz.ch

## ABSTRACT

Internet background radiation (IBR) is a very interesting piece of Internet traffic as it is the result of attacks and misconfigurations. Previous work has primarily analyzed IBR traffic to large unused IP address blocks called network telescopes. In this work, we build new techniques for monitoring one-way traffic in live networks with the main goals of 1) expanding our understanding of this interesting type of traffic towards live networks as well as of 2) making it useful for detecting network service reachability problems. Our first contribution is a comprehensive classification scheme for dissecting one-way traffic into useful classes, including one-way traffic due to unreachable services, scanning, peer-to-peer applications, and backscatter. Our classification scheme is helpful for monitoring IBR traffic in live networks solely based on flow-level data. After validating our classifier, we use it to analyze a massive data-set that covers 7.41 petabytes of traffic from a large backbone network to shed light into the composition of one-way traffic. We report a number of interesting observations including that one-way traffic makes a very large fraction, i.e., between 34% and 67%, of the total number of flows to the monitored network, although it only accounts for only 3.4% of the number of packets, which suggests a conceptual model for Internet traffic in which IBR is dominant in terms of flows. In addition, we find that the vast majority of one-way traffic is due to scanning. Finally, we demonstrate the utility of one-way traffic of the particularly interesting class of unreachable services for monitoring network services by describing interesting outages we detected in the network of our university.

## 1. INTRODUCTION

Studying the background radiation of the Internet has been instrumental for understanding Internet threats, like the prevalence of Denial of Service (DoS) attacks [23] and the propagation of Slammer [21] and Blaster [3] worms. IBR traffic has been primarily studied with the help of network telescopes [22], which are large unpopulated IP address blocks (e.g. /8) that seamlessly filtering out all benign traffic. In live networks, IBR traffic can be extracted by dissecting two-way from one-way traffic, i.e., traffic connections that do not receive a network reply. In this work, *we build techniques for monitoring one-way traffic in live networks.*

We make three main contributions. First, we *design*

*and validate a comprehensive classification scheme for dissecting one-way traffic* that relies solely on flow-level data, e.g., NetFlow. Our scheme uses 17 carefully-designed rules, which encode novel heuristics as well as leverage and combine existing techniques that have been proven effective for certain types of traffic. We associate each one-way flow with up to 18 different *signs*, which help dig into the causes of one-way flows. The key features of our scheme are that: 1) it relies on flow data; 2) it is very easy to configure; 3) it provides transparency into the classification process; and 4) it is easily extensible.

Our second contribution is that *we use our classification scheme to shed light into IBR traffic towards a large live network by analyzing a massive dataset of unsampled flow records summarizing 7.41 petabytes of traffic* recorded over eight years (from 2004 to 2011). We find that in the studied years one-way traffic makes between 34% and 67% of the total number of flows, but it only accounts on average for 3.4% of the packets. This suggests a conceptual model of Internet traffic in which IBR traffic is a needle in a haystack in terms of packets (and bytes), but it is very dominant in terms of flows. Besides, we find that scanning is the largest component of one-way traffic corresponding to 83.5% and 62.6% of the one-way flows and packets, respectively. The second major contributor is traffic caused by P2P applications. Finally, we find that in 2011 the volume of IBR traffic in terms of flows is almost equal to 2004 with small fluctuations in the years between. The fraction of IBR traffic in the total number of flows has been consistently decreasing from 67% in 2004 to 34% in 2011.

Third, *we show how one-way flows assigned to the class of unreachable services provide new possibilities for service availability monitoring enabling in particular to passively assess the impact of outages*, which is not possible with traditional approaches based on active probing. To demonstrate the utility of our scheme, we describe and measure the impact of interesting outages and misconfiguration we discovered in the network of our university.

In the next section we provide preliminary insights into the one-way traffic classification problem. Next, in Section 3 we describe our datasets and their sanitization steps. Our classification scheme is introduced in

Section 4 and validated in Section 5. In Section 6 we use our scheme to characterize the composition of IBR traffic in a live network, while in Section 7 we outline how our work is useful for service reachability monitoring. Finally, we present related work and conclude our paper in Section 8 and 9.

## 2. PRELIMINARIES

Network communication inherently reflects dialogs taking place between applications running on distributed systems including any accompanied control and error messages. Regular communication involves two-way traffic consisting of packets in opposite direction. In contrast, one-way traffic results from communication errors or other unusual situations that result in packets in one direction. In this work, we look at one-way communication at the network level. We define a *two-way flow* (or equivalently a *bi-flow*) as the aggregate of two standard 5-tuple flows that have reverse values in the source and destination IP address and port number fields. A 5-tuple traffic flow is a flow of packets that have common values in the IP address, port number, and protocol fields. We define an *one-way flow* as a flow that does not have a matching reverse flow. A two-way flow may still indicate a failed communication at the transport or application level, e.g., a TCP reset packet produces a two-way flow, therefore *one-way flows provide a conservative view of failed communication attempts* in the Internet.

One-way traffic is important as it is associated with interesting events, such as unreachable services, scanning, and misconfigurations. To draw useful information from one-way traffic we need to infer what is causing it. The goal of one-way traffic classification is to classify one-way traffic into classes of malicious and benign causes. We identify three main causes of one-way traffic:

- **Failures & Policies:** this group comprises *attempts to access a service that do not succeed* due to failures or policies. Possible reasons are firewall blocking, temporary or permanent outages of network elements, stopped service processes on end-hosts, and routing misconfigurations. For example, the use of laptops outside of enterprise networks they are configured for produces one-way traffic. Similarly, one-way traffic is the result of attempts to use services that are blocked by security policies, like when testing connectivity with the ping tool to non-responding destinations. Furthermore, one-way traffic can result from network congestion.
- **Attacks:** *vulnerability scanning* is a very common attack vector. However, unsolicited probes typically do not receive a reply resulting in many one-way flows. Another attack case is *backscatter traffic* from DoS attacks using spoofed source IP addresses, which also results in one-way flows towards the spoofed IP addresses. Third, *prefix*

*hijacking* results in traffic towards a “black hole”.

- **Special application behavior:** *P2P applications* try to reconnect to systems they have been in contact before as such systems are often preferred peers by application design. However, the dynamic nature of P2P networks and in particular nodes unexpectedly leaving a network result in one-way reconnection attempts. A second, but rare scenario is *applications using sessions that run the two communication directions over different connections*. For example, an interesting case we have identified is the application layer protocol of LANsource [12] that runs over UDP and uses different destination ports for each communication direction.

Properly identifying Internet one-way traffic from measurement data imposes specific requirements on monitoring. In particular, traffic sampling and poor spatial network coverage, i.e., missing one direction of a communication that may take a different path due to asymmetric routing, may result in artificial one-way flows. In principle, dealing with asymmetric routing requires to either know how flows are routed or to monitor every link along a network cut between a monitored network and the Internet. For these reasons, one-way traffic can in practice be measured near the edge, e.g., in enterprise, university, or small ISP networks, where it is possible to monitor the entire border traffic activity without sampling. Edge networks are the vast majority of Internet domains.

## 3. DATASETS AND SANITIZATION

We use data from a regional academic backbone network that serves 46 single-homed universities and research institutes. The monitored address range contains 2.2 million IP addresses, which correspond to a continuous block slightly larger than a /11. We have been archiving unsampled flow records since 2003 from hardware-based NetFlow meters on the border routers of the monitored network. The meters capture all the traffic that crosses the border destined to or coming from the Internet. In a single peering link, we observe in 2011 on average 108.1 million flows per hour, which corresponds to 3,064 million packets. The stored flow records are not anonymized. From each record we extract for our classification scheme the following fields: IP addresses, port numbers, protocol number, byte/packet counts, and timestamps. We do not use TCP flags because they are not supported by the fast hardware-based NetFlow, although they can be easily integrated in our classification. Over time we have accumulated a massive archive with more than 100 TBytes of compressed NetFlow records.

A major challenge in analyzing our data is computational overhead especially for tasks that cannot be parallelized. For example, using an optimized C++ program to parse one day of data and compute the average number of one-way flows per host takes 10 hours and 10-20

Gbytes of main memory in a dedicated high-end cluster. The computational overhead is dominated by the time for decompression and for disk I/O. To make the computational overhead more tractable, we use two 400-hour samples per year, which correspond to approximately one month per year. In particular, we analyze the first 400 hours of each February and August between 2004 and 2011. The studied traffic data correspond to 457 gigafloes and 7.41 petabytes and cover approximately 9% of the total number of flows observed during the eight year period. This is one of the largest traffic traces that has been analyzed in the literature.

For each 400-hour sample, in the top rows of Table 1 we report the time the monitoring and collection infrastructure was functioning properly and the traffic volume it corresponds to. Overall the availability of the infrastructure was higher than 99.5% with only short interruptions due to router restarts and corrupted data. To eliminate the impact of interruptions in our analysis, we use daily averages and discard the days in which the short interruptions occurred.

### 3.1 Data Sanitization

In this section we describe how we addressed a number of data sanitization challenges for extracting one-way flows. In particular, we outline our methodology for eliminating double-counting, defragmenting flows, and pairing bi-flows. To speed up searching for matching flows, we split time into intervals and limit our search within an interval. In Section 3.1.4 we analyze the impact of the interval size. In the future, the standardization of bidirectional flow monitoring [27] will likely make one-way flow extraction much easier.

#### 3.1.1 Double-Counting Elimination

Although we collect flows from border routers only, if a flow is routed through multiple border routers, it will be reported multiple times in our NetFlow data. To address this problem we use a map provided to us by SWITCH network operators that specifies which ports of the border routers are connected to the internal backbone and which to the Internet. Our NetFlow data carry information about the input and output port of a flow. We use this information to filter out flows. We only keep flows that are routed between the Internet and the intranet. Double-counting elimination has a significant impact on our data reducing the total traffic volume by 32.3%.

#### 3.1.2 Defragmentation

We call the standard 5-tuple NetFlow records *raw flows*. NetFlow may fragment a traffic flow when the active timeout expires or when the flow cache is filled up and flushed. In addition, the optional NetFlow *fast aging* feature, with which our hardware based meters are configured, leads to additional fragmentation. Fast flow aging reduces the utilization of the flow table by exporting a flow if no more than a few packets are observed within the first few seconds of the creation of a flow table entry. It realizes an efficient process to age

out entries for short flows. To address flow fragmentation due to fast aging, we split time into intervals and for each raw flow we search within its time interval for other raw flows with the same 5-tuple. We merge matching flows updating their byte/packet counts and start/end timestamps. We call the merged flows *defragmented flows*. Note that if a flow spans multiple intervals it will yield one defragmented flow for each interval. In the bi-flow pairing process we account for these cases. In the third and fourth row of Table 1 we show the numbers of raw and defragmented flows. We observe that defragmentation has a significant impact reducing the number of flows by a fraction ranging between 20.6% and 39.6% for different years.

#### 3.1.3 Bi-flow Pairing

For TCP and UDP a two-way flow is the aggregate of two defragmented flows that have the same 5-tuple with reverse values in the source and destination IP address and port fields. For each defragmented flow in an observation interval, we search for flows in the same or in the adjacent intervals and group matched flows into a bi-flow. An one-way flow is a flow that does not have a matching reverse flow in the same or in an adjacent interval.

A special situation exists for other protocols, like ICMP, Encapsulated IPv6, and Encapsulated Secure Payloads (ESP). The port fields in our NetFlow data do not carry interesting semantics for protocols other than TCP and UDP. For these protocols we define a flow and match bi-flows based on a 3-tuple consisting of the IP addresses and protocol fields. This definition may underestimate the number of one-way flows involving protocols other than TCP and UDP.


Interestingly, it is not officially documented how NetFlow defines a flow for these protocols. According to certain sources NetFlow aggregates packets with common values in the source/destination IP address, protocol, and ICMP type/code fields into an ICMP flow [29]. For protocols other than TCP, UDP, and ICMP, NetFlow sets port fields to zero [29]. We find that this behavior is not met by our NetFlow data. Only 5% of all ICMP flows carry the ICMP type and code information, while all remaining ICMP flows have this information set to zero.

#### 3.1.4 Impact of the Interval Size

We next analyze the impact of the interval size by computing flow metrics for a full day of data using different interval sizes. We select a 10-minute reference interval within which we can search efficiently for flow fragments and pairs. In Figure 1 we show the impact of the interval size with respect to the 10-minute (600 secs) interval. We observe, as expected, that longer intervals lead to slightly lower flow counts as fewer flows cross interval boundaries. This decrease is less for two-way flows due to their longer duration. We observe that doubling the interval size slightly decreases absolute count metrics by 3-5%. On the other hand, the

	2004	2005	2006	2007	2008	2009	2010	2011
Hours	799.0	800.0	792.0	780.0	799.0	799.0	800.0	800.0
Total terabytes	270.8	304.5	424.4	641.5	842.6	1035.2	1763.2	2132.7
Raw flows (1e9)	30.96	29.35	30.32	49.03	63.16	77.05	90.91	86.49
Defragmented flows (1e9)	24.31	22.47	22.53	29.62	44.03	58.55	72.21	66.91
Two-way flows (1e9)	6.117	6.712	8.046	10.84	17.61	22.97	28.71	26.55
Two-way packets (1e9)	381.1	418.6	558.0	834.0	1080.8	1286.8	2092.3	2372.7
One-way flows (1e9)	12.19	9.102	6.502	8.040	8.967	12.78	15.01	14.06
One-way packets (1e9)	34.09	20.89	24.04	31.09	27.59	49.65	39.56	79.03

**Table 1: Size of data sets per year in total hours, bytes, flows and packets. The different flow counts represent the processing steps from raw data to defragmented flows and finally to flows separated into two- and one-way flows.**



**Figure 1: Impact of time interval size on flow metrics aggregated for a full day. The decrease of the fraction of one-way to total flows is not more than 1.2% for a doubled interval size.**

decrease of the relative volume of one-way flows is only 1.2% and does not decrease further with an increasing interval size. Therefore, we conclude that with a 10-minute interval absolute count metrics slightly overestimate the number of flows, while relative count metrics, which is the type of metrics we primarily report, are very accurate.

## 4. ONE-WAY TRAFFIC CLASSIFICATION

Our classification scheme has the following key features:

- **Flow-based Classification:** We leverage solely flow monitoring data, i.e., NetFlow. Compared to packet-level data and DPI that have been used in the past to monitor IBR traffic, flow monitoring provides fewer details, but is less expensive, scales better to monitoring large IP address blocks, and is more generally available.
- **Easy Configuration:** Our classifier does not require training, which is a drawback of many traffic classification techniques. It can be applied with minimal configuration, without extensive overhead.
- **Comprehensibility:** We use legible rules instead of complex classification structures, like self-organizing maps, to make our classifier comprehensible to non-machine-learning experts. Our scheme is based

on rules summarized in Table 3 that can be easily reviewed and verified both by researchers and network operators. Traffic classification schemes should be easily understandable to be engaged and deployed by network operators.

- **Extensibility:** Our classifier can be easily extended with additional traffic features and classification rules. For example, if a flow monitor is configured to collect additional data, like TCP flags, then these can be easily integrated into new or existing classification rules.

We classify one-way flows using a two step process that resembles the way medical diagnosis identifies the cause of a health problem. We first define a set of *signs*, which are informative attributes for digging into the cause of one-way flows. We then check each flow against all signs and find the ones it matches. Finally, we classify a flow based on the collection of signs that are present or absent. Our classification scheme is based on a carefully designed set of rules that take as input the signs of a flow and determine the appropriate class, including the unknown class.

### 4.1 Signs


Flow signs are derived from *sign tests*. A sign test determines if a flow exhibits a specific sign and may be as simple as checking a flow feature value or as complex as performing a behavioral assessment of an involved end-host. We introduce in total 18 signs exploiting in 4 cases techniques from the literature that have been shown effective for detecting scan and peer-to-peer traffic. We classify our signs in four categories based on the entity they characterize: 1) a pair of communicating hosts, 2) a remote host as a source of an one-way flow, 3) a local host as a target of an one-way flow, or 4) a flow. In Table 2 we summarize our signs and show in which category they belong.

#### 4.1.1 Host-Pair Signs

Host-pair signs characterize a pair of end-hosts with an one-way flow. They reflect the mix of one- and two-way flows of a host pair. In Figure 2 the graphs a to d illustrate communication patterns that correspond to the four possible combinations of a biflow and an outflow between a host pair with an inflow. Hosts are represented by nodes, where the local host is the node in the

Sign Type	Sign Name	Detection Criterion/Algorithm	Sign Short Name for Rules
Host pair behavior	End-hosts-communicating Limited dialog	One-way flow between productive host pair One-way flows between unproductive host pair	PotOK InOut
Remote host behavior	Service sole reply Remote scanner 1 Remote scanner 2 Remote non-scanner	no biflow on $\text{srcIP} \wedge \text{dstPort} \geq 1024 \wedge \text{srcPort} < 1024$ TRW algorithm (suspected scanner) Host classification (suspected scanner) TRW algorithm (suspected regular host)	Backsc TRWscan HCscan TRWnom
Local host behavior	Unused local address Service unreachable Peer-to-peer	Unpopulated local IP address Unanswered request to local service Flow towards local P2P host	GreyIP Unreach P2P
Flow feature	Artifact Single packet Large flow Bogon Protocol Retry	UDP/TCP flow with both port numbers=0 Flow contains one packet only Flow carries $\geq 10$ packets or $\geq 10240$ bytes Source IP belongs to bogon space IP protocol type of flow Average inter-packet arrival time $> 256$ sec	Artef Onepkt Large Bogon TCP, UDP, ICMP, OTHER Retry

**Table 2: Overview of defined signs based on the behavior of one or both involved hosts or based on the numeric values of one or more flow attributes.**



**Figure 2: Mixture of incoming one- and two-way flows exchanged between a host pair shown as a graph. Hosts are represented by nodes and the presence of inflow/outflow/biflows by arrows.**

left, and flows with arrows. We use the sign *end-hosts-communicating* to mark one-way flows between hosts that are involved in a productive communication, i.e., they have a biflow between them, with occasional failed connections, i.e., one-way flows. We assign this sign to one-way flows with the communication pattern shown in the graphs b and d in Figure 2.

Furthermore, we introduce a sign called “InOut” to capture the pattern of graph c in Figure 2. In particular, the “InOut” sign identifies host pair communication situations in which an one-way ICMP flow is a reply to another outgoing one-way flow exchanged between the same host pair. We assign ICMP one-way flows with the sign “InOut” to a special class that includes suspected benign one-way flows.

#### 4.1.2 Remote Host Signs

We use four signs to characterize remote hosts that generate one-way flows towards the monitored network. If we observe one-way flows from a well-known source port of a remote host to a local ephemeral port and at the same time we do not observe any two-way flows nor opposite direction one-way flows between the two hosts, then we assign the *service sole reply* sign. We use this sign to detect backscatter from attack traffic towards a service of the remote host. We check for the absence of biflows within a 30-minute time interval to limit computational overhead.

Besides, we leverage two state-of-the-art algorithms to detect remote scanners. We introduce three signs that describe either a positive or a negative test result. The first two signs are based on the well-known Thresh-

old Random Walk (TRW) [16] algorithm, which uses sequential hypothesis testing to qualify a remote host as scanner or benign (or undecidable). We denote these signs with “TRWscan” and “TRWnom”. The second algorithm [2] is a variant of TRW that inspects the behavior of a remote host in a relaxed way and classifies it as a scanner if a test condition is met. A negative test result is not supported, i.e., it does not qualify a host as benign. The test assesses the *service fanout* of a remote host, i.e., the count of 2-tuples  $\{\text{dstIP}, \text{dstPort}\}$  the host attempted to access, and its overall success in connection establishment. Each connection attempt is qualified either as “good” if it results in a two-way flow or as “bad” if it results in an one-way flow. This way, a remote host is classified as a scanner if it has a bad service fanout larger than two times the good service fanout (and at least 4). The host classification is done in a first pass. During a second pass any flow sourced by a classified scanner is labeled with the sign “HCscan”. Both algorithms restrict their tests to 30 minute time windows to keep memory demands in bounds.

#### 4.1.3 Local Host Signs

We use three signs to characterize local hosts within the monitored network. We mark with the sign *unused local address* (“GreyIP”) local IP addresses that never sourced any flow over a full observation period of 400 hours and, therefore, are unpopulated.

We introduce the sign *service unreachable* to trace outages of well-known services. We define a service as an endpoint described by the 3-tuple  $\{\text{localIP}, \text{protocol}, \text{localPort}\}$  with the constraint that the port number is well-known, i.e., below 1024. We consider only TCP and UDP flows. We consider a local service valid if it serves at least 20 clients during any 30-minute interval over a full observation period. We label one-way flows that target valid services with the sign *service unreachable* (“Unreach”).

The sign *peer-to-peer* marks one-way flows caused by stale host caches of P2P applications. Behavioral P2P application identification has been studied extensively [17, 25, 15]. We make use of three rules from the literature (H1,2,3 from [15]) that identify a local host as a member of a P2P network based on the observation of typical

P2P behavioral patterns. We require that at least 2 out of 3 P2P rules is matched or a match occurs for at least 5 time windows of 30 minute duration. First, we prepare a list with detected local P2P hosts by analyzing two-way flows of a full observation period. Then, we mark any TCP/UDP one-way flow targeted at a high port towards such hosts with the *peer-to-peer* sign.

#### 4.1.4 Flow Signs

The remaining signs are either direct flow attributes or are derived from them by the rules shown in Table 2. The sign *artifact* is used for TCP and UDP flows with both port numbers set to zero. Such flows are the result of packet fragments for which the correct port numbers are not available due to the missing layer-4 header. If a flow contains a source address that belongs to bogon space using the data of [8] then we mark it with the sign *bogon*. We use the sign *retry* to address the situation where a host re-uses the same 5-tuple for a connection retry by canceling the TCP TIME\_WAIT state through non-standard configuration of the socket interface.

## 4.2 Classifier

A class is associated with one or more rules. We define the following classes:

- **Service Unreachable:** access attempt to temporary unavailable services.
- **Malicious Scanning:** probing for the exploitation of vulnerabilities in end systems.
- **Benign P2P:** P2P applications trying to access peers in their local host cache that are not anymore available.
- **Backscatter:** replies to DoS attack traffic that uses randomly chosen source IP addresses to hide the real identity of an attacker.
- **Suspected Benign:** one-way flows may exist as part of benign applications using data and control connections in parallel and employing one of them for acknowledgment only. Another cause may be temporary failures within an otherwise productive communication.
- **Bogon:** one-way flows originating from bogon IP space.
- **Other:** one-way flows that do not match any of the above classes.

We carefully design our classification rules with a systematic process shown in Figure 3. We start with an initial set of rules and in each iteration we classify flows, compute a conflict report of the resulting classification and update the rules. The conflict report describes flows classified in multiple classes. Specifically, it reports: 1) the classes that intersect; 2) the size of the intersections; 3) a list of sign combinations in the intersections ordered by their popularity; and 4) the number of remaining unclassified flows. Based on the conflict report

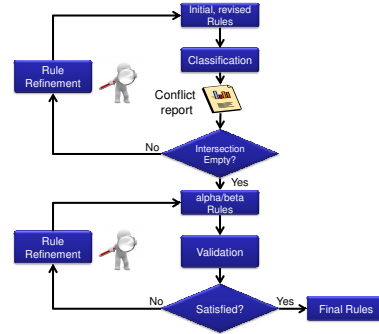


Figure 3: Rule refinement stages

and in particular the sign combinations, we then manually update the rules to resolve conflicts and to reduce the number of unclassified flows. Each flow is checked against 18 different signs, which in theory yields a maximum of 262,144 possible sign combinations. However, most combinations do not occur in practice. In our data, we observe a total of 1,035 different sign combinations. Moreover, one could ignore sign combinations that are only observed in a small number of flows resulting in a sharp reduction on the number of interesting combinations. We iteratively repeat this procedure until we resolve all conflicts and cannot further significantly reduce the number of unclassified flows. Based on the described procedure we derived a first version of our classification rules after eight iterations. Then, we further refine our rules based on our validation. The final classifier includes 17 classification rules shown in Table 3.

The classes “Malicious Scanning” and “Suspected Benign” employ multiple rules with different contribution to the overall class results. Based on our data, we see two rules that contribute most to the class “Malicious Scanning”. In particular, rules 3 and 4 match 73.5% and 10.2% of the flows in this class, respectively. The high value of rule 4 is surprising as this rule contains two seemingly conflicting signs (TRWnom and HCscan). This indicates that the TRW algorithm rates some sources too optimistic as being nominal. The remaining rules account for 9: 5.8%, 7: 5.4%, 6: 2.0%, 8: 1.6%, 2: 0.9%, 5: 0.5% and 1: 0.3%. In the class “Suspected Benign” the major contributor is rule 15 with 56.9%, followed by rule 13 with 24.0%, rule 16 with 16.7% and rule 14 with 2.4%. We discuss the accuracy of our ruleset in Section 5.

Our classifier processes input flows in three passes. In a first pass it initializes local services, local P2P hosts and inactive local IP addresses. In a second pass matching signs are assigned to flows and in a final pass rules are applied to account each flow to a unique class. We have implemented our classifier in C++ and run it on a dedicated Linux cluster with a fast RAID disk subsystem. We collect logs of the overhead of the three passes combined for all 16 runs. In general, the over-

Class Name	Rule #	Flow Membership Rules
Malicious Scanning	1	$\{\overline{TRWscan}, \overline{HCscan}, \overline{PotOk}\} \Rightarrow Scanner$
	2	$\{\overline{HCscan}, \overline{TRWscan}, \overline{TRWnom}, \overline{PotOk}\} \Rightarrow Scanner$
	3	$\{\overline{TRWscan}, \overline{HCscan}, \overline{PotOk}\} \Rightarrow Scanner$
	4	$\{\overline{TRWnom}, \overline{HCscan}\} \Rightarrow Scanner$
	5	$\{\overline{GreyIP}, \overline{Onepkt}, \overline{TRWscan}, \overline{HCscan}, \overline{Backsc}, \overline{ICMP}, \overline{UDP}, \overline{bogon}\} \Rightarrow Scanner$
	6	$\{\overline{GreyIP}, \overline{TRWscan}, \overline{HCscan}, \overline{Onepkt}, \overline{ICMP}, \overline{Backsc}, \overline{bogon}\} \Rightarrow Scanner$
	7	$\{\overline{Onepkt}, \overline{GreyIP}, \overline{ICMP}, \overline{TRWscan}, \overline{HCscan}, \overline{TRWnom}, \overline{bogon}, \overline{P2P}, \overline{Unreach}, \overline{PotOk}, \overline{Backsc}, \overline{Large}\} \Rightarrow Scanner$
	8	$\{\overline{GreyIP}, \overline{Onepkt}, \overline{TRWscan}, \overline{HCscan}, \overline{Backsc}, \overline{ICMP}, \overline{TCP}, \overline{bogon}\} \Rightarrow Scanner$
	9	$\{\overline{ICMP}, \overline{TRWscan}, \overline{TRWnom}, \overline{HCscan}, \overline{InOut}, \overline{bogon}, \overline{PotOk}\} \Rightarrow Scanner$
Backscatter	10	$\{\overline{Backsc}, \overline{TRWscan}, \overline{HCscan}, \overline{P2P}, \overline{InOut}, \overline{PotOk}\} \Rightarrow Backscatter$
Service Unreachable	11	$\{\overline{Unreach}, \overline{TRWscan}, \overline{HCscan}, \overline{bogon}, \overline{P2P}\} \Rightarrow Unreachable$
Benign P2P Scanning	12	$\{\overline{P2P}, \overline{TRWscan}, \overline{HCscan}, \overline{bogon}\} \Rightarrow P2P$
Suspected Benign	13	$\{\overline{PotOk}, \overline{Unreach}, \overline{P2P}, \overline{TRWnom}, \overline{bogon}\} \Rightarrow Benign$
	14	$\{\overline{Large}, \overline{GreyIP}, \overline{TRWscan}, \overline{HCscan}, \overline{P2P}, \overline{Unreach}, \overline{PotOk}, \overline{ICMP}, \overline{Backsc}, \overline{bogon}, \overline{TRWnom}\} \Rightarrow Benign$
	15	$\{\overline{TRWnom}, \overline{GreyIP}, \overline{HCscan}, \overline{P2P}, \overline{Unreach}, \overline{bogon}, \overline{Backsc}\} \Rightarrow Benign$
	16	$\{\overline{ICMP}, \overline{InOut}, \overline{TRWscan}, \overline{HCscan}, \overline{TRWnom}, \overline{bogon}, \overline{PotOk}\} \Rightarrow Benign$
Bogon	17	$\{\overline{bogon}, \overline{TRWscan}, \overline{HCscan}, \overline{Backsc}\} \Rightarrow Bogon$

**Table 3: Rules used to classify one-way flows. Each rule specifies which signs have to be present or absent (sign names with overbars). An overview of defined signs can be found in Table 2.**

head scales with the size of the dataset, but also depends on dataset characteristics. On average a run processed 28.6 gigafloWS in 23.1 hours (with a maximum of 37.4 hours) and required 3.4 GB of memory (with a peak of 6.1 GB). Considering the monitoring period length of 400 hours these figures are well within real-time bounds. Note that we exclude the overhead for preprocessing (time binning, flow defragmentation, biflow-pairing) as it largely depends on the available input data (e.g. NetFlow export or YAF-created flow records). On our NetFlow records the preprocessing required the most resources while still safely remaining within real-time bounds.

Our classification scheme has one parameter that needs configuration: the time window size used for P2P and services identification and for scanner detection. We set it by default to 30 minutes, but on a smaller network the window size should be increased (see also Section 5).

## 5. VALIDATION

In this section we evaluate the accuracy of our classifier and optimize its rule set. Validating an one-way traffic classification scheme is very challenging. A major challenge is the scarcity of information available from one-way traffic consisting of very short flows carrying few packets and frequently no payload. We address this challenge, first, by building more accurate host profiles based on flows over long time windows, second, by carefully examining one-way packets for violation of relevant protocol state machines, and third, by exploiting DPI and application identification techniques.

We first use a traffic summarization technique [5] based on frequent item-set mining (FIM) to summarize flows of different classes into frequent item-sets, which can be then easily inspected by an analyst. We apply this approach in each class and measurement period between 2004 and 2011. In summary, we find that

for different classes and periods up to 75% of the flows were correctly classified. However, with this approach we cannot assess the remaining flows and the number of false negatives.

Realizing a thorough evaluation solely with flow data is challenging, as we miss information like TCP flag and ICMP types/codes. A validation of one-way traffic classification solely based on flow data provides insufficient evidence. For this reason, we built a dedicated validation setup in a smaller network to obtain more detailed data.

### 5.1 Validation Setup

We built a monitoring setup and collected packet traces at the Internet gateway of the swiss university Hochschule für Technik Rapperswil (HSR) occupying a /16 IP range. The campus network is sparsely populated with a total of 3,949 active IP addresses seen during our experiment. We use a dedicated host that is tapping traffic between the border router and the firewall. On the firewall no filtering is activated, enabling us to see all traffic routed to this IP range. To reliably collect packet data, we use an Endace DAG card that provides exact packet timestamps. Mandated by the IT security policy, we anonymize all IP addresses using a prefix-preserving scheme. Note that anonymization does not allow us to use active probing for our validation as proved to be useful in [24].

We collected packet traces for 19 consecutive days and extracted flows using the YAF flow meter [13] configured for bidirectional flow export. We then split flows into 10-minute intervals and defragmented flows within each interval resulting in a total of 322.7 mio flows, which include 219.6 mio (68.1%) incoming one-way flows, 8.06 mio outgoing one-way flows (2.5%) and 95.1 mio two-way flows (29.5%). The small fraction of two-way flows can be explained by the scarce popula-

tion of the monitored network. We retain packet data for a sub-period of three randomly chosen days.

To obtain application labels we configured YAF with its optional DPI application identification feature that assigns application labels to flows it recognizes. However, the coverage achieved by YAF is limited. Thus, we used an additional DPI application identifier to add a second set of application labels for almost all flows. This additional DPI application identifier originally was developed to evaluate the BLINC classifier [18] and in an improved version to evaluate several competing approaches to application identification [19].

To learn more on the occurrence of one-way flows using bogus source IP addresses we periodically downloaded the full bogon list provided by [8] throughout the data collection process. This list not only contains the official IANA-reserved address ranges, but includes all IP ranges assigned to regional registrars that have not yet been handed out to costumers.

## 5.2 Validation Criterias

To determine class memberships we build a DPI classifier that uses an extensive set of 34 rules that make heavy use of details available only from packet-level data. In the following paragraphs, we summarize the key additional information we extract from the packet-level data and how we exploit it. Then, we provide more details on detecting the different classes.

**Extended Host Profiles:** For each host, we maintain an extended host profile over a full observation period of 457 hours (19 days) that tracks successful and failed connections. We use extended host profiles in the TRW algorithm and, in combination with application identification, for identifying P2P hosts and local services. This is feasible in a /16 network, i.e., it required 12 GBytes of memory for 10.2 mio extended host profiles, but does not scale reasonably for larger networks.

**ICMP types and codes:** For ICMP flows, we analyze if type and code information and the communication situation of the involved host pair fit a class. For example, we deem a flow as backscatter if it is an echo reply and the receiver never sent a request to the sender. In addition, we check if the source of an ICMP flow is not already identified as a scanner, which interestingly is true for more than 92.1% of all incoming one-way ICMP flows.

**Protocol State Machine:** As an additional way to identify malicious scan traffic, we analyze how well individual flows follow the transport layer protocol state machine. Malformed packets are frequently used to exploit weaknesses of protocol stacks or to penetrate non-stateful firewalls. For TCP flows, we test if the flags of consecutive packets fit into acceptable state changes in the TCP protocol machine and run sanity checks on the sequence numbers of segments.

**Application Identification:** We apply application identification techniques [13, 18, 19] on two-way flows to discover local services and to detect hosts running P2P applications.

**Precise Timestamps:** We rely on packets timestamps to identify the initiator of a connection, which is not possible using NetFlow timestamps. Initiator detection is useful in combination with DPI and extended host profiles for mapping local services.

The rules used for validation and the corresponding DPI-based signs are described in Tables 4 and 5, respectively. In the next few paragraphs, we provide a summary of the rules of each class.

### 5.2.1 Class ‘Malicious Scanning’

For a scanner detection, we rely on the TRW algorithm which we configure for an upper bound on false positives of 0.5% and a IP population degree of 20

We counted a total of 10.2 mio unique IP addresses outside of our network during the 457 hours: 8.50 mio are pure sources of one-way flows, 1.22 mio are involved in two-way flows only, and 0.446 mio external hosts have both one- and two-way flows. Dissecting the group having one-way flows only we get 7.62 mio external hosts involved in inflows, 0.628 mio involved in outflows and a remaining group of 0.254 mio hosts involved in both types of one-way flows. Finally, we observe an almost linear growth of the number of unique external IP addresses over time, showing only a small decrease of the growth rate after two-third of the observation period. This indicates that scanner frequently change their IP address and/or use a high count of randomly chosen spoofed IP addresses for a decoy scanning that hides the real source behind faked source addresses.

As an additional way to identify malicious scan traffic we analyze how well individual flows follow the transport layer protocol state machine by inspecting arriving packets. This criterion is based on the fact that malformed packets frequently are used to exploit weaknesses of protocol stacks or to penetrate non-stateful firewalls. For TCP flows we test if the flags of consecutive packets fit into acceptable state changes in the TCP protocol machine and run sanity checks on the sequence numbers of segments. For UDP flows we cannot gain much insight from the protocol state machine. Thus, we just test if a packet has no payload to detect simple UDP scans. In the case of ICMP traffic, we inspect the ICMP type and code information picked from the packet header. Depending on the type information, we can distinguish between ICMP query and error messages. But, it is hard to draw conclusions from this information as almost any kind of ICMP messages is popular for scanning including ICMP destination port unreachable messages [10]. We classify unanswered ICMP requests as scan flows without further checks. For other one-way ICMP flows, we evaluate for each flow the communication situation of the host pair involved (e.g., is this ICMP flow an answer to a UDP flow?) and also check if the flow source is not already identified as a scanner, which is true for more than 92.1% of all incoming one-way ICMP flows. Otherwise, if we see that there are solely inflows exchanged between a host pair, then we classify an ICMP flow as backscatter.



Sign Type	Sign Name	Detection Criterion/Algorithm	Sign Short Name for Rules
Host pair behavior	End-hosts-communicating Limited dialog	One-way flow between productive host pair One-way flows between unproductive host pair	PotOK InOut
Remote host behavior	Service sole reply Remote scanner Remote nominal Peer-to-peer source	No biflow on $\text{srcIP} \wedge \text{dstPort} \geq 1024 \wedge \text{srcPort} < 1024$ (no SYN flag) TRW algorithm (suspected scanner) TRW algorithm (suspected nominal) UDP/TCP flow sourced by a known P2P host	BackscPort ScanSrc NoScan P2pSrc
Local host behavior	Unused local address Temporary Inactive Service unreachable High Service unreachable Peer-to-peer destination	Unpopulated local IP address Populated but currently inactive IP address Unanswered request to local service Unanswered request to local service on high port UDP/TCP flow targeted at a known P2P host	DarkIP OffDst SrvDst HiSvDst P2pDst
Flow feature	P2P flow ICMP error 1 ICMP error 2 ICMP request ICMP reply TCP flooding ICMP other UDP/TCP scan flow Reserved IP Bogon IP	P2P flow (by DPI application identification) ICMP error message 'port unreachable' ICMP error message 'unreachable' (non-port) ICMP request message ICMP reply message TCP flow with > 50 packets, all with same flags Remaining ICMP messages not fitting any of the other ICMP signs TCP flow violating state machine or UDP flow wit no payload Source IP is a IANA reserved address Source IP is from full bogon IP address list (unallocated)	P2pApp IcmpErr1 IcmpErr2 IcmpReq IcmpReq TcpFlood IcmpOth ScanInv BogonIana BogonFull

**Table 4: Overview of defined validation signs based on the behavior of one or both involved hosts, DPI application identification or based on numeric values of one or more flow or packet attributes.**

Finally, we count all remaining incoming TCP/UDP one-way flows towards unpopulated IP addresses as scan traffic unless they already match a more specific rule of another class. This class assignment follows the principle of Occam’s razor.

### 5.2.2 Class ‘Backscatter’

Backscatter traffic is created either as a reaction to DoS attacks using spoofed source IP addresses or is the result of routers sending ICMP error messages, e.g. of type ‘network unreachable’. Furthermore, any answer to an ICMP request containing a spoofed source IP address can be considered backscatter traffic. For each of these types we test separately. We detect DoS-initiated backscatter traffic when we see a TCP or UDP flow with a low source port ( $< 1024$ ) and a high destination port ( $\geq 1024$ ) sourced by a host that never receives traffic from our network during the full observation period (and for TCP does not contain SYN packets). We consider ICMP unreachable error messages as backscatter whenever they are not sourced by an identified scanner and there are solely inflows exchanged between the host pair involved. Moreover, ICMP replies are identified as backscatter traffic whenever they are sourced by an external IP address that never receives traffic from our network and at the same time was not identified as a scanner.

### 5.2.3 Class ‘Service Unreachable’

We describe a service offered by our network as a three-tuple  $\{localIP, protocol, localPort\}$ . To elect a service we count the external sources of two-way flows that are identified by DPI inspection to be of a well-known application type and that are initiated by the external host. For initiator-detection we rely on precise packets timestamps provided by our test setup. If the count of unique external IP addresses sourcing such flows exceeds a threshold of 20 then we add a service to the services list. We build the services list along the host profiles during a first pass over the full dataset. This way we detected 36 local services.

### 5.2.4 Class ‘Benign P2P Scanning’

We use DPI application identification applied to two-way flows to identify hosts running productive P2P applications. To qualify a host as a P2P user we set a minimum threshold of 10 on the number of identified P2P two-way flows we see during the full observation period. Thus, we base the P2P host detection on the host behavior profiles we extended by P2P flow counts. We identified a total of 26,393 external and 410 local P2P hosts. We cannot detect Skype traffic because its proprietary protocol significantly changed in its current version making evaluated detection rules unusable. This underestimates the number of flows falling into this class.

### 5.2.5 Class ‘Suspected Benign’

This class takes all one-way flows that do not seem to be malicious and are not yet assigned to one of the other benign classes. Examples are flows sourced by misconfigured applications that target unpopulated IP space or failed connection attempts between a host pair, which has a bi-flow. To become a candidate, a flow must not be sourced by an identified scanner. This includes flows adhering to the TCP protocol state machine or being sourced by a host identified to be nominal. In addition, we consider repeated valid TCP connection retries or any flows sourced by P2P hosts towards populated but currently inactive IP addresses as potentially benign. A last group is formed by ICMP error messages and replies sent to local hosts that sourced an outgoing flow to the external IP address they receive the ICMP message from.

### 5.2.6 Class ‘Bogon’

Flows falling into this class use source IP addresses from IP ranges not allocated to any organization. Therefore, such flows are likely the result of IP address spoofing in the context of decoy scanning or direct denial-of-service attacks.

To identify all IP addresses belonging to unallocated

Class Name	Rule #	Flow Membership Rules
Malicious Scanning	1	$\{\overline{ScanInv}, \overline{ScanSrc}, \overline{NoScan}, \overline{P2pApp}, \overline{P2pSrc}, \overline{P2pDst}, \overline{SrvDst}, \overline{HiSvDst}, \overline{PotOk}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Scanner$
	2	$\{\overline{ScanSrc}, \overline{ScanInv}, \overline{IcmpErr1}, \overline{IcmpErr2}, \overline{IcmpReq}, \overline{IcmpRpl}, \overline{P2pApp}, \overline{P2pSrc}, \overline{P2pDst}, \overline{HiSvDst}, \overline{BogonFull}\} \Rightarrow Scanner$
	3	$\{\overline{ScanInv}, \overline{ScanSrc}, \overline{P2pSrc}, \overline{P2pApp}, \overline{DarkIP}, \overline{SrvDst}, \overline{HiSvDst}, \overline{BogonFull}\} \Rightarrow Scanner$
	4	$\{\overline{TcpFlood}, \overline{OffDst}, \overline{ScanSrc}, \overline{NoScan}, \overline{P2pDst}\} \Rightarrow Scanner$
	5	$\{\overline{ScanInv}, \overline{ScanSrc}, \overline{DarkIP}, \overline{BackscPort}, \overline{BogonFull}\} \Rightarrow Scanner$
	6	$\{\overline{IcmpErr1}, \overline{ScanSrc}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Scanner$
	7	$\{\overline{IcmpErr2}, \overline{ScanSrc}, \overline{InOut}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Scanner$
	8	$\{\overline{IcmpReq}, \overline{P2pDst}, \overline{PotOk}, \overline{BogonIana}, \overline{BogonFull}, \overline{NoScan}\} \Rightarrow Scanner$
	9	$\{\overline{IcmpRpl}, \overline{ScanSrc}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Scanner$
	10	$\{\overline{DarkIP}, \overline{OffDst}, \overline{ScanInv}, \overline{ScanSrc}, \overline{IcmpRpl}, \overline{NoScan}, \overline{BackscPort}, \overline{IcmpReq}, \overline{BogonFull}, \overline{IcmpErr1}, \overline{IcmpErr2}\} \Rightarrow Scanner$
Backscatter	11	$\{\overline{BackscPort}, \overline{P2pSrc}, \overline{P2pDst}, \overline{ScanSrc}, \overline{PotOk}, \overline{BogonFull}, \overline{InOut}\} \Rightarrow Backscatter$
	12	$\{\overline{IcmpRpl}, \overline{ScanSrc}, \overline{InOut}, \overline{PotOk}, \overline{P2pDst}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Backscatter$
	13	$\{\overline{IcmpErr1}, \overline{ScanSrc}, \overline{InOut}, \overline{PotOk}, \overline{P2pDst}, \overline{BogonIana}, \overline{BogonFull}, \overline{NoScan}\} \Rightarrow Backscatter$
	14	$\{\overline{IcmpErr2}, \overline{ScanSrc}, \overline{InOut}, \overline{PotOk}, \overline{P2pDst}, \overline{BogonIana}, \overline{BogonFull}, \overline{NoScan}\} \Rightarrow Backscatter$
Service Unreachable	15	$\{\overline{SrvDst}, \overline{P2pDst}, \overline{ScanSrc}, \overline{TcpFlood}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Unreachable$
	16	$\{\overline{SrvDst}, \overline{P2pDst}, \overline{ScanInv}, \overline{ScanSrc}, \overline{P2pSrc}, \overline{NoScan}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Unreachable$
	17	$\{\overline{SrvDst}, \overline{P2pDst}, \overline{ScanInv}, \overline{ScanSrc}, \overline{P2pSrc}, \overline{DarkIP}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Unreachable$
Benign P2P Scanning	18	$\{\overline{P2pApp}, \overline{P2pSrc}, \overline{P2pDst}\} \Rightarrow P2P$
	19	$\{\overline{P2pSrc}, \overline{P2pDst}, \overline{P2pApp}\} \Rightarrow P2P$
	20	$\{\overline{P2pDst}, \overline{OffDst}, \overline{P2pSrc}, \overline{ScanSrc}, \overline{DarkIP}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow P2P$
	21	$\{\overline{P2pDst}, \overline{P2pApp}, \overline{OffDst}, \overline{ScanSrc}, \overline{P2pSrc}, \overline{DarkIP}, \overline{SrvDst}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow P2P$
	22	$\{\overline{HiSvDst}, \overline{P2pDst}, \overline{ScanSrc}, \overline{TcpFlood}\} \Rightarrow P2P$
	23	$\{\overline{P2pApp}, \overline{P2pDst}, \overline{P2pSrc}, \overline{OffDst}, \overline{BogonIana}, \overline{BogonFull}\} \Rightarrow P2P$
	24	$\{\overline{P2pApp}, \overline{P2pSrc}, \overline{P2pDst}, \overline{ScanSrc}, \overline{DarkIP}, \overline{BackscPort}, \overline{TcpFlood}, \overline{BogonIana}, \overline{BogonFull}, \overline{PotOk}\} \Rightarrow P2P$
Suspected Benign	25	$\{\overline{NoScan}, \overline{BackscPort}, \overline{P2pSrc}, \overline{P2pDst}, \overline{SrvDst}, \overline{BogonIana}, \overline{PotOk}\} \Rightarrow Benign$
	26	$\{\overline{TcpFlood}, \overline{OffDst}, \overline{ScanSrc}, \overline{P2pApp}, \overline{NoScan}, \overline{P2pDst}, \overline{DarkIP}\} \Rightarrow Benign$
	27	$\{\overline{P2pSrc}, \overline{OffDst}, \overline{P2pDst}, \overline{TcpFlood}, \overline{DarkIP}, \overline{ScanSrc}\} \Rightarrow Benign$
	28	$\{\overline{PotOk}, \overline{ScanSrc}, \overline{P2pDst}, \overline{SrvDst}\} \Rightarrow Benign$
	29	$\{\overline{IcmpErr1}, \overline{InOut}, \overline{ScanSrc}, \overline{DarkIP}, \overline{OffDst}, \overline{NoScan}, \overline{PotOk}\} \Rightarrow Benign$
	30	$\{\overline{IcmpRpl}, \overline{InOut}, \overline{ScanSrc}, \overline{DarkIP}, \overline{OffDst}, \overline{NoScan}, \overline{PotOk}\} \Rightarrow Benign$
	31	$\{\overline{IcmpErr2}, \overline{InOut}, \overline{ScanSrc}, \overline{DarkIP}, \overline{OffDst}, \overline{NoScan}, \overline{PotOk}\} \Rightarrow Benign$
Bogon	32	$\{\overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Bogon$
	33	$\{\overline{BogonFull}, \overline{BogonIana}\} \Rightarrow Bogon$
	34	$\{\overline{BogonIana}, \overline{BogonFull}\} \Rightarrow Bogon$

**Table 5: Rules used to classify one-way flows based on packet-level data. Each rule specifies which signs have to be present or absent (sign names with overbars). An overview of defined signs can be found in Table 4.**

IP space, we use an extended list beyond the IANA reserved IP ranges that we download from a periodically updated database [8].

### 5.2.7 Class ‘Other’

All flows which do not match any of the defined rules are assigned to class “Other”. Typically this includes flows carrying no or few signs or that, in rare cases, carry contradictory signs excluding them from any classification rule.

## 5.3 Validation Results

We run both the flow-based classifier and our DPI classifier over the evaluation dataset. Then, we compare the results of both classification runs and fill-in a *confusion matrix* showing predicted (flow-based) versus actual (DPI-based) class memberships. Based on the confusion matrix, we calculate the *recall* (also called sensitivity), *specificity* and *precision* metrics for all defined classes.

For the flow-based classification we use the same settings as we apply to the full dataset with the exception of the time window parameter for remote host profiling, which we increase by a factor of 32 to 16 hours instead of 30 mins. This compensates for the smaller

probability to observe randomly targeted scan flows in the evaluation dataset caused by the smaller IP address range. Similarly to the full-dataset analysis, we classify incoming one-way flows. Our validation uses the full evaluation dataset to build host profiles and a subset of 34.8 mio incoming one-way flows for DPI. For this subset we could validate 99.4% of all one-way flows.

The test results are summarized in Table 6. We obtain very good results for the classes “Malicious Scanning”, “Service Unreachable” and “Benign P2P Scanning” with recalls above 95.3%, specificity values greater than 91.3% and precision exceeding 95.4%. The class “Suspected Benign” has a recall of 85.1% and a precision of 75% caused by flows in the borderline between the classes “Suspected Benign” and “Malicious Scanning” that are hard to separate. On the other hand, its specificity is excellent at 99.9%. For the class “Backscatter” we observe a recall of 62.4%. We miss primarily ICMP replies and error messages as ICMP type and code information is not available in our flow data. However, the recall can be easily improved by using ICMP type/code fields when available to find ICMP replies. Besides, a specificity of 100% and a precision of 88.4% shows this class is still very useful. Finally, we observe a low recall for the class “Bogon” of 40.4% and

Class Name	Recall [%]	Specificity [%]	Precision [%]
Malicious Scanning	99.9 (99.6)	91.2 (88.8)	99.8 (99.6)
Backscatter	62.4 (62.4)	100.0 (100.0)	88.4 (87.2)
Service Unreachable	99.6 (99.5)	100.0 (100.0)	96.1 (92.5)
Benign P2P Scanning	95.3 (91.8)	99.9 (99.8)	95.5 (85.0)
Suspected Benign	85.1 (70.5)	99.9 (99.9)	75.0 (70.5)
Bogon	40.4 (40.2)	100.0 (100.0)	100.0 (100.0)

**Table 6: Results of validation.** The values in parentheses provide a lower bound by counting non-validated flows as classification mismatches.

at the same time a perfect specificity and precision of 100.0%. This is because the bogon list used by the flow-based classifier is a subset of the bogon list used by our DPI classifier. We intentionally exclude the full-bogon information from flow-based classification because historical full-bogon data is not available. Again, it is trivial to extend the flow-based classifier to use full-bogon lists for analyzing present or future data.

## 5.4 Impact on Flow Classifier


For flow-based scan detection we use two sets of rules. The first set is based on scan detection algorithms providing strong evidence. The second set defines rules that point towards scanning, but provide weaker evidence. In particular, it includes flows 1) towards unpopulated IP addresses or 2) that consist only of a single packet and do not match any other rule. The validation shows that the second rule set in fact is very useful for scan detection, which led us to add this rule set to the class “Malicious Scanning” eliminating an initially defined class “Suspicious Other”. In particular, 65.1% of the flows matching these rules are actually detected as scanners when observing host behavior for an extended time period.

Another change introduced during validation concerns the assignment of ICMP flows. We observe that 92.1% of all ICMP flows seen are sourced by identified scanners. Without ICMP type/code information, it is difficult to satisfactorily distinguish backscatter ICMP flows from scanning. Consequently, we moved a rule initially assigned to the class “Backscatter”, assuming ICMP replies from DoS attacks, to the class “Malicious Scanning”. Furthermore, we added the sign “InOut” that helps to identify host pair communication situations in which an ICMP flow is a reply to an one-way flow exchanged between the same host pair. We use this sign in the class “Suspected Benign”.

## 6. ONE-WAY TRAFFIC COMPOSITION

In this section we apply our classification scheme on a massive dataset of flows records collected between 2004 and 2011 to shed light into the composition and characteristics of one-way traffic.

**Aggregate Statistics:** We first find that in terms of flows one-way traffic is a very large component of Internet traffic. During the studied period one-way flows correspond to between 34% and 67% of the total number of flows towards the monitored network. This is important for systems that need to keep per flow state,



**Figure 4: Composition of one-way traffic in flow counts per class.**

like stateful firewalls and flow meters. On the other hand, due to their short-lived nature, one-way traffic corresponds only to 3.4% and 0.79% of the total number of packets and bytes, respectively. This indicates that IBR traffic is not a significant problem in terms of bandwidth consumption and packet overhead.

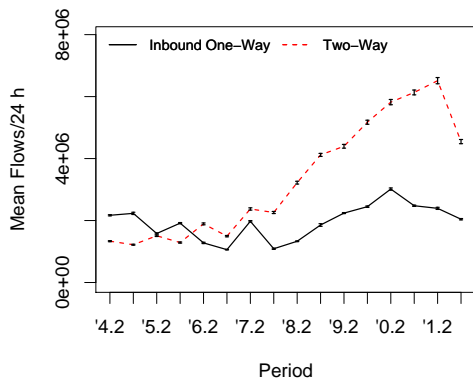
Figure 4 shows how one-way traffic breaks down into different classes in terms of flows. In addition, Table 7 provides for each measurement period and class the fraction of one-way traffic in terms of flows and packets. On average the class “Malicious Scanning” accounts for 83.5% of all one-way flows. It is followed by the classes “Benign P2P” (6.7%), “Service Unreachable” (4.8%), “Suspected Benign” (2.6%), “Other” (2.2%), “Backscatter” (0.3%) and “Bogon” (0.1%), where in parentheses we show the average fraction of one-way flows per class over the eight year period. The packet perspective provides a different view. “Malicious Scanning” accounts for 62.6% of all one-way packets, “Benign P2P” for 13.0%, “Service Unreachable” for 10.1%, “Suspected Benign” for 9.1%, “Other” for 4.7%, “Backscatter” for 0.5%, and “Bogon” for 0.03%.

We observe that one-way traffic is clearly dominated by scanning. In terms of packets, scanning accounts for a smaller fraction of 62.6% of the total number of one-way packets. This is because one-way flows classified as scanning consist on average of 1.6 packets per flow. In contrast, one-way flows of the classes “Service Unreachable” and “Benign P2P” consist of 4.1 and 12.1 packets per flow, respectively. This difference is because TCP one-way flows to unreachable services are much more persistent in SYN packet retransmission attempts than TCP scanning. In addition, we find that one-way flows in the class “Benign P2P” are often multi-packet UDP flows that sharply increase the average number of packets per flow.

**Changes Over Time:** Besides, we observe a number of notable changes over time. In Figure 5 we compare how the mean daily number of one- and two-way flows evolved between 2004 and 2011. The volume of

Period	Malicious Scanning	Backscatter	Unreachable	Benign P2P	Suspected Benign	Bogon	Other
2004-02	87.2%/57.9%	0.1%/0.2%	4.8%/9.2%	5.5%/10.3%	1.1%/3.7%	0.1%/0.1%	1.1%/18.7%
2004-08	93.9%/82.0%	0.0%/0.1%	0.6%/1.3%	3.4%/9.8%	1.0%/4.1%	0.0%/0.0%	1.0%/2.7%
2005-02	87.2%/65.6%	0.0%/0.1%	5.8%/12.1%	4.1%/9.1%	1.8%/10.5%	0.0%/0.0%	1.0%/2.7%
2005-08	95.4%/74.7%	0.0%/0.1%	0.4%/0.7%	1.8%/4.3%	1.5%/5.3%	0.3%/0.1%	0.7%/14.9%
2006-02	89.3%/71.8%	0.1%/0.2%	0.6%/4.4%	4.2%/9.9%	3.3%/10.2%	0.2%/0.1%	2.3%/3.5%
2006-08	87.8%/47.3%	0.1%/0.0%	0.4%/0.3%	5.3%/4.1%	3.2%/46.9%	0.0%/0.0%	3.1%/1.4%
2007-02	86.9%/83.6%	0.1%/0.1%	1.7%/1.2%	6.5%/7.1%	2.8%/6.2%	0.0%/0.0%	2.0%/1.8%
2007-08	80.7%/63.5%	0.2%/0.2%	2.2%/2.9%	9.5%/17.6%	3.5%/10.3%	0.2%/0.1%	3.7%/5.5%
2008-02	68.3%/34.4%	0.2%/0.2%	1.1%/1.9%	23.9%/49.4%	3.0%/10.4%	0.0%/0.0%	3.5%/3.7%
2008-08	70.5%/47.8%	0.3%/0.4%	15.2%/26.3%	6.5%/11.4%	3.2%/8.5%	0.0%/0.0%	4.4%/5.5%
2009-02	74.2%/55.7%	0.1%/0.0%	15.3%/25.3%	5.9%/11.1%	2.7%/5.6%	0.0%/0.0%	1.8%/2.3%
2009-08	86.2%/75.2%	0.1%/0.1%	3.7%/6.0%	5.7%/10.2%	2.3%/5.2%	0.0%/0.0%	2.0%/3.3%
2010-02	88.2%/77.5%	0.3%/0.2%	2.6%/5.1%	4.0%/7.7%	2.9%/5.9%	0.1%/0.0%	2.0%/3.5%
2010-08	85.8%/74.3%	0.2%/0.2%	3.1%/4.7%	6.3%/12.5%	2.6%/5.5%	0.0%/0.0%	1.9%/2.7%
2011-02	81.4%/61.5%	0.6%/1.2%	6.2%/12.9%	7.3%/18.7%	2.6%/3.5%	0.0%/0.0%	1.9%/2.3%
2011-08	72.8%/28.3%	2.2%/4.3%	12.3%/46.7%	6.9%/15.1%	3.7%/4.3%	0.0%/0.0%	2.0%/1.4%

**Table 7: Fraction of flows/packets falling into the defined one-way flow classes. The class *Other* represents the remainder of flows not captured by the other classes. Note, that we do not list rules that achieve a low coverage eliminating the flow artifacts class that results from fragmented packets without layer-4 header.**



**Figure 5: Evolution of one and two-way mean daily flow counts over time. The vertical bars mark 95% confidence intervals.**

one-way flows has exhibited small fluctuations between 2004 and 2011 and, interestingly, in 2011 it is almost equal to 2004. On the other hand, the number of two-way flows has grown significantly by 343%. A study by Akamai estimated the global Internet penetration increase in 2010 as 17% [1]. If we extrapolate this figure over eight years, it yields a growth by 351%, which is very close to our observations. Besides, we observe that the fraction of one-way flows has dropped over time. In 2004 one-way flows accounted for 67% of the total number of flows. Their share gradually dropped and since 2007 one-way flows account consistently for one out of three incoming flows.

The fraction of “Malicious Scanning” flows varies, started at a high of 87.2% in 2004 and stayed at this or a higher level until 2006. Then, in 2007 we note a decline that leads to a low of 68.3% in Aug. 2008. This decline correlates with an increase of the class P2P one-way traffic that falls early into the same period and also an increase of class “Unreachable” later on. The high counts in class “Service Unreachable” between Aug. 2008 and Feb. 2009 are caused by unan-

swered NTP requests towards a NTP server. Investigating further this incident, we found that the increased number of requests pushed the operator to introduce access restrictions to this server. As we analyze more in the next section, this class of one-way flows is very useful to administrators to monitor the reachability of services. The peak in benign P2P one-way traffic between Aug. 2007 and Feb. 2008 can be attributed to two IP addresses that are reported to be fake servers by the eMule Bulletin Board. Fake eMule servers are commonly used to collect eMule user data for research or legal prosecution. Minor contributors to this class are XBSlink traffic sourced by Xbox 360 and PS2/3 game equipment, and the MMORPG (massively multiplayer online role-playing game) Heroes of Might and Magic.

## 7. SERVICE REACHABILITY MONITORING

In this section, we leverage one-way flows of the class “Service Unreachable” to demonstrate a new approach to monitor local services and detect reachability problems. Traditionally, the availability of local network services within an enterprise, a data center, or a university network is monitored using active probing tools or server logs. In contrast, we exploit traffic flow data and our classification scheme to monitor local services from the network. Our approach has the following key advantages: 1) *It provides a tangible assessment of the impact of disruptions* by enabling to count the actual number of remote client IP addresses that fail to reach a service; 2) *It exploits passive measurements* in contrast to injecting network overhead for active probing; 3) *It automatically discovers running local services* circumventing the need to manually configure (new) services; 4) *It enables to concurrently monitor many services* by taking advantage of data from the strategical network location of gateways.

In the next paragraphs, we first outline our analysis methodology and secondly we discuss interesting misconfigurations and outages we detected in the campus

network of ETH Zurich by applying our scheme on one week of data.

## 7.1 Methodology

Using one-way flows classified with our scheme in the class “Service Unreachable” as well as bidirectional flows as input, we compute the following three metrics to evaluate the reachability of local services within the network of ETH Zurich. We focus on services running on well-known port numbers.

We define the *availability* of a service as the number of time intervals a service is alive, i.e., it has bidirectional flows, over the number of intervals it is either alive or unresponsive, i.e., it has only incoming one-way flows. We ignore intervals in which a service does not have any incoming flows. The availability metric helps group services into permanently and temporarily offered ones.

Available services are not necessarily reachable to *all* clients due to access policies, misconfigurations, or other reachability problems. We define the *reachability* of a service as the number of distinct client IP addresses with bidirectional flows to the service over the total number of distinct client IP addresses to the service during a time interval. The reachability metric is useful for identifying misconfigurations and outages.

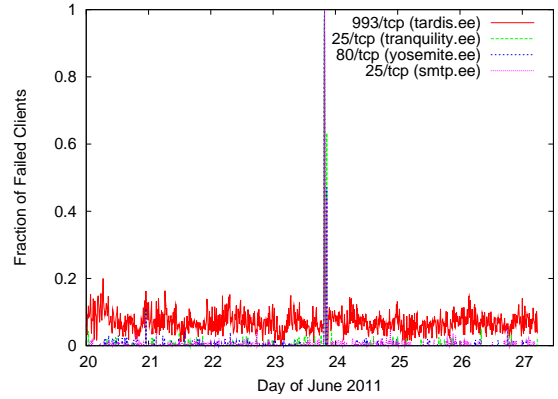
Finally, the *outage impact* of a service is the number of unique client IP addresses that are involved in one-way flows of the class “Service Unreachable” during a time interval. As several clients can be behind a remote NAT, this metric provides a lower bound on the actual number of clients that are affected by an outage and is useful for assessing the impact of failures.

## 7.2 Outages and Misconfigurations

In this section we demonstrate the utility of one-way flows detected with our classification scheme for analyzing service reachability problems.

We applied our scheme on a randomly-selected week of NetFlow data from our archive and initially focused on services discovered within the address range of the Department of Information Technology and Electrical Engineering of ETH Zurich. We discovered services as 2-tuples  $\{localIP, localPort\}$  with a well-known port number that had more than 20 unique clients during any time interval. As with bi-flow pairing, we used a 10-minute time interval. We discovered in total 43 remotely accessible departmental services for which we computed our three metrics. To validate our observations, we met and discussed with the network administrators of our department providing them details about the availability, reachability, and outages of the services. The administrators confirmed the validity of the discovered services and provided feedback on a number of interesting outages we summarize in the next paragraphs.

We identified a group of 32 services available at least 99% of the time and 11 services that either fail very often or are not operated all the time. For all frequently used services we found a coinciding global outage on the 23rd of June 2011 starting at approximately 19:40 UTC



**Figure 6: Coinciding outage observed on the 23rd of June 2011 19:40 UTC on most university services resulting in failed accesses from 287,583 unique clients.**

and lasting more than 10 minutes as shown in Figure 6. This was a general problem in the reachability of the department network from the outside. To learn more about this outage, we analyzed three key services of the entire university: the main web, webmail, and software distribution services of ETH Zurich. Surprisingly, we found that at the same point in time the services were not accessible indicating a campus-wide temporary reachability problem. During the identified 10-minute interval the outage impact metric revealed that 287,583 unique clients failed to access target services! Our investigation showed that the outage occurred during a planned router software upgrade. Although, temporal disruptions due to the upgrade were expected, the scale of the disruption revealed was surprisingly larger than expected. Present server availability monitoring techniques based on active probing and server logs miss this important information. In addition, we visually inspected the time series of the volume of one-way flows for the top departmental services and found that the outage resulted in a sharp increase in requests targeting mail services due to the automated retry behavior of mail clients. On the other hand, for web services the volume of one-way flows during the outage did not substantially differ from the volume of two-way flows during regular operation, indicating a more network-friendly retry behavior.

Our inspection of the reachability metric for the discovered services revealed two actionable cases. First, we found an NTP server that was consistently not reachable on average by 12.9% of its clients. The outage impact metric showed that in this case too a surprisingly large number of clients failed to receive a reply: during the studied week in total 2.2 million unique clients were involved in one-way flows to the NTP server. Investigating this further unveiled that standard NTP “hello world” configuration examples use the `swisstime.ee.ethz.ch` server in our department, making it likely one of the most popular NTP servers in the Internet. In addition, `swisstime.ee.ethz.ch` is often preconfigured

in embedded devices. The failure rate, i.e., fraction of one-way flows, varies for different countries of origin. In the top client countries we observe Belgium with failure rate 65.7%, India with 20.2%, and Germany with 16.8%. Our investigation found two main causes of low reachability: 1) an IP filter was configured to prevent access from a /8 address block, which was frequently the source of attacks; and 2) certain devices appear to compute invalid CRC checksums resulting in dropped requests. From these findings we learned both that the traffic flow volume to the studied network was heavily affected from the popular NTP server and that a misconfigured application computes invalid CRC checksums.

Second, we noticed that the SMB service offered on 445/tcp was targeted by far more one-way flows (112% more) than two-way flows. An in-depth analysis revealed that almost all one-way flows were sourced by a likely misconfigured client that persistently tried to connect to this service. All remaining services achieved a reachability of 97.1% or more with most of them ranking at the 99.9% figure.

## 8. RELATED WORK

Studying IBR traffic has been very useful for understanding Internet threats. In 2001, Moore *et al.* [23] were the first to use network telescopes to monitor background radiation and answer questions such as “how prevalent are denial-of-service attacks in the Internet today?”. Yegneswaran *et al.* [31] employed scan logs collected by Dshield [9] during four months in 2001 and 2002 to analyze quantitative characteristics of intrusion activity in the global Internet. They made a number of interesting observations including that a very small collection of correlated sources are responsible for a significant fraction of intrusion attempts. In 2004, Pang *et al.* [24] used traces, lasting 80-hours to one week, from three sites to conduct a detailed characterization of traffic towards four large unused address blocks. Among other observations, they highlighted the dynamism of background radiation, which unlike normal traffic, can exhibit significant changes on a nearly daily basis. More recently, Wustrow *et al.* [30] examined background radiation between 2006 and 2010 to four /8 unused blocks. They observed among other findings that background radiation continues to be very dynamic and identified interesting artifacts that lead to address space pollution.

A number of more recent studies have focused on one-way traffic. In 2012, Brownlee [6] described a tool called *iatmon* for classifying one-way traffic based on packet interarrival times and remote scanner patterns of the schema of Treurinet [28]. In addition, the author described observations about the different types of scanning in six months of data from the UCSD Network Telescope [7]. Our work is complementary to the *iatmon* tool, as we focus on flow instead of packet-based classification. Besides, the work by Lee and Brownlee [20] used three one-day packet traces captured in

2000, 2003, and 2006 in the University of Auckland to analyze the lifetime and size distribution of one and two-way flows. In our work, we provide new insights into one-way flows based on a substantially larger trace and introduce novel methods to classify their causes. Allman *et al.* [2] used TCP connection logs from the Lawrence Berkeley National Laboratory collected between 1994 and 2006 to characterize how scanning changed over time. In addition, they correlated observed patterns to well-known malware activity. Our study follows on until 2011 and focus primarily on how to classify one-way traffic. Jin *et al.* [14] used flow data towards grey IP space of a campus network and developed a heuristic to identify external scanners. We follow along this direction and highlight the need to classify one-way traffic and identify a number of different causes in addition to scanning.

Finally, a number of studies have focused on similar but less related problems. Guha *et al.* [11] characterized the fraction of useful flows generated by mobile hosts in an enterprise network and found that 34% of the flows fail mainly due to hosts mobility. Schatzmann *et al.* [26] described a flow-based system to identify reachability problems in remote networks caused by events like prefix hijacking. Barford *et al.* [4] characterized the network location of sources of background radiation and derived a model for generating source IP addresses in simulation and emulation environments with realistic distributional properties.

## 9. CONCLUSIONS

One-way traffic towards live networks is an exotic piece of Internet traffic that accounts for a very large fraction of the total number of flows. To shed light into this type of traffic, we design and validate a classification scheme for one-way traffic. Classifying one-way traffic based solely on flow-level data is very challenging as one-way flows contain very few packets often without any payload. To address these challenges, we design and carefully optimize 18 classification rules that leverage primarily communication patterns of involved end hosts.

We use our classifier to analyze the composition of a massive dataset of traffic flow records from a regional backbone network and make the following key observations:

- 1) One-way traffic makes a very large fraction of the total number of flows, although it only accounts for only 3.4% of the number of packets
- 2) Scanning accounts for 83.5% and 62.6% of the one-way flows and packets; peer-to-peer applications of 6.7% of the flows and 13.0% of the packets; and unreachable services for 4.8% of the flows and 10.1% of the packets.
- 3) The number of two-way flows has grown significantly by 343% between 2004 and 2011. In contrast, the absolute volume of one-way flows in 2011 is almost equal to 2004.
- 4) The share of one-way flows has declined from 67% in 2004 to 34% in 2011.

Finally, we focus on the particularly interesting class of one-way traffic caused by unreachable services and show how this class can be very useful for monitoring the reachability of network services and in particular for assessing the impact of outages.

In the past IBR traffic has been primarily studied with the help of a small number of large network telescopes. Our techniques are useful for network operators of edge networks for monitoring their local IBR traffic and for detecting service reachability problems. For this reason, with this paper we will make the code of our classification scheme publicly available. In addition, to shed further light into the composition of one-way traffic and to facilitate correlating data from multiple IBR traffic monitors, we will release non-sensitive details about the evolution of one-way and two-way traffic between 2004 and 2011 based on the studied datasets. In particular, for each 10-minute time interval, we will make available statistics about the synthesis of one-way traffic, the popularity of different signs, the targeted services, and the geographical distribution of sources.

## Acknowledgements

We are grateful to SWITCH for providing us precious data for studying one-way traffic. In addition, we are thankful to Dominik Schatzmann for helping with double-counting elimination and to Simon Leinen for providing us valuable feedback on service reachability monitoring. Lastly, we wish to acknowledge anonymous reviewers for their helpful comments and suggestions.

## 10. REFERENCES

- [1] Akamai Technologies. The state of the internet report (3rd quarter, 2009). Technical report, 2009.
- [2] M. Allman, V. Paxson, and J. Terrell. A brief history of scanning. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, page 82. ACM, 2007.
- [3] M. Bailey, E. Cooke, F. Jahanian, D. Watson, and J. Nazario. The blaster worm: Then and now. *IEEE Security and Privacy*, 3:26–31, July 2005.
- [4] P. Barford, R. Nowak, R. Willett, and V. Yegneswaran. Toward a model for source addresses of internet background radiation. In *Proc. of the Passive and Active Measurement Conference*. Citeseer, 2006.
- [5] D. Brauckhoff, X. Dimitropoulos, A. Wagner, and K. Salamati. Anomaly extraction in backbone networks using association rules. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, IMC '09, pages 28–34, New York, NY, USA, 2009. ACM.
- [6] N. Brownlee. One-way traffic monitoring with iatmon. In *Passive and Active Measurement Conference*, 2012.
- [7] CAIDA. UCSD Network Telescope. <http://www.caida.org/data/realtime/telescope/>.
- [8] T. Cymru. The Bogon Reference. <http://www.team-cymru.org/Services/Bogons/>, 2012.
- [9] Cooperative Network Security Community - Internet Security. [www.dshield.org](http://www.dshield.org).
- [10] F. Gont. ICMP attacks against TCP. RFC 5927, July 2010.
- [11] S. Guha, J. Chandrashekar, N. Taft, and K. Papagiannaki. How healthy are today's enterprise networks? In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, pages 145–150. ACM, 2008.
- [12] IANA - Internet Assigned Numbers Authority. PORT NUMBERS. <http://www.iana.org/assignments/port-numbers>, 2011.
- [13] C. Inacio and B. Trammell. Yaf: yet another flowmeter. In *Proceedings of the 24th international conference on Large installation system administration*, pages 1–16. USENIX Association, 2010.
- [14] Y. Jin, Z. Zhang, K. Xu, F. Cao, and S. Sahu. Identifying and tracking suspicious activities through IP gray space analysis. In *Proceedings of the 3rd annual ACM workshop on Mining network data*, page 12. ACM, 2007.
- [15] W. John and S. Tafvelin. Heuristics to classify internet backbone traffic based on connection patterns. *International Conference on Information Networking (ICOIN)*, pages 1–5, 2008.
- [16] J. Jung, V. Paxson, A. Berger, and H. Balakrishnan. Fast portscan detection using sequential hypothesis testing. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 211–225, 2004.
- [17] T. Karagiannis, A. Broido, and M. Faloutsos. Transport layer identification of p2p traffic. *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, pages 121–134, 2004.
- [18] T. Karagiannis, K. Papagiannaki, and M. Faloutsos. Blinc: multilevel traffic classification in the dark. In *ACM SIGCOMM Computer Communication Review*, volume 35, pages 229–240. ACM, 2005.
- [19] H. Kim, K. Claffy, M. Fomenkov, D. Barman, M. Faloutsos, and K. Lee. Internet traffic classification demystified: myths, caveats, and the best practices. In *Proceedings of the 2008 ACM CoNEXT conference*, page 11. ACM, 2008.
- [20] D. Lee and N. Brownlee. Passive measurement of one-way and two-way flow lifetimes. *SIGCOMM Comput. Commun. Rev.*, 37(3):17–28, 2007.
- [21] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the slammer worm. *IEEE Security and Privacy*, 1:33–39, July 2003.
- [22] D. Moore, C. Shannon, G. Voelker, and S. Savage. Network telescopes: Technical report. Technical report, 2004.
- [23] D. Moore, G. M. Voelker, and S. Savage. Inferring internet denial-of-service activity. In *SSYM'01*:

- Proceedings of the 10th conference on USENIX Security Symposium*, Berkeley, USA, 2001.
- [24] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson. Characteristics of Internet background radiation. In *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, pages 27–40. ACM New York, NY, USA, 2004.
- [25] M. Pernyi, T. D. Dang, A. Gefferth, and S. Molnr. Identification and analysis of peer-to-peer traffic. *JOURNAL OF COMMUNICATIONS*, 1(7), 2006.
- [26] D. Schatzmann, S. Leinen, J. Kogel, and W. Muhlbauer. Fact: Flow-based approach for connectivity tracking. In *Passive and Active Measurement conference*, Mar. 2011.
- [27] B. Trammell and E. Boschi. Bidirectional flow export using IPFIX. RFC 5103, January 2008.
- [28] J. Treurniet. A network activity classification schema and its application to scan detection. *IEEE/ACM Trans. Netw.*, 19(5):1396–1404, Oct. 2011.
- [29] Wikipedia. Netflow. <http://en.wikipedia.org/wiki/Netflow>.
- [30] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Huston. Internet background radiation revisited. In *Proceedings of the 10th annual conference on Internet measurement*, IMC '10, pages 62–74, New York, NY, USA, 2010. ACM.
- [31] V. Yegneswaran, P. Barford, and J. Ullrich. Internet intrusions: global characteristics and prevalence. *SIGMETRICS Perform. Eval. Rev.*, 31(1):138–147, 2003.