# Hershel: Single-Packet OS Fingerprinting

Zain Shamsi, Ankur Nandwani, Derek Leonard, *Student Member, IEEE*, and Dmitri Loguinov, *Senior Member, IEEE*

***Abstract*—Traditional TCP/IP fingerprinting tools (e.g., nmap) are poorly suited for Internet-wide use due to the large amount of traffic and intrusive nature of the probes. This can be overcome by approaches that rely on a single SYN packet to elicit a vector of features from the remote server. However, these methods face difficult classification problems due to the high volatility of the features and severely limited amounts of information contained therein. Since these techniques have not been studied before, we first pioneer stochastic theory of single-packet OS fingerprinting, build a database of 116 OSs, design a classifier based on our models, evaluate its accuracy in simulations, and then perform OS classification of 37.8M hosts from an Internet-wide scan.**

***Index Terms*—Internet measurement, OS fingerprinting.**

## I. Introduction

With the explosive growth and distributed nature of computer networks, it has become progressively more difficult to manage, secure, and identify Internet devices. One of the approaches that greatly helps in understanding the composition of internal and external networks is *OS fingerprinting*, which is a process that determines the operating system of remote hosts based on peculiarities of their network-level behavior. This differentiation is possible due to certain freedom in selection of default stack parameters, ambiguities in IETF RFCs [8], [35], [36], noncompliant TCP/IP implementations, and lacking standardization for responses to malformed requests.

Over the last 20 years, these observations have led to a variety of methods [3]–[5], [7], [9], [17], [23], [27], [44]–[46], [49], [50], [53], [55]–[57] that perform classification using application-layer traffic, TCP/IP/UDP headers, ICMP packets, or some combination thereof. These algorithms are useful not only in network security (i.e., detection of outdated/unpatched hosts), but also market analysis [30] and Internet characterization [18], [24], [31], [34], [49]. However, their usage, scalability, and accuracy at very large scale (i.e., millions of destinations) have not been explored before. We aim to address this issue in the following.

Manuscript received July 16, 2014; revised March 06, 2015; accepted June 15, 2015; approved by IEEE/ACM Transactions on Networking Editor S. Sen. Date of publication July 06, 2015; date of current version August 16, 2016. The work of D. Loguinov was supported by the NSF under Grants CNS-1017766 and CNS-1319984. An earlier version of the paper appeared in ACM SIGMETRICS 2014.

Z. Shamsi and D. Loguinov are with Texas A&M University, College Station, TX 77843 USA (e-mail: zain@cs.tamu.edu; dmitri@cs.tamu.edu).

A. Nandwani is with Coinbase, Inc., San Fransisco, CA 94104 USA (e-mail: ankur.nandwani@gmail.com).

D. Leonard is with Acxiom Corporation, Little Rock, AR 72201 USA (e-mail: dleonard@cse.tamu.edu).

Digital Object Identifier 10.1109/TNET.2015.2447492

### A. Methodology and Objectives

The Internet has been the target of numerous measurement studies, with the trend recently shifting from covering a small subset of destinations [34], [39] to scanning the entire IP space [12], [18], [24], [40]. This allows researchers to enumerate live hosts, detect vulnerabilities, and shed light on deployment of new protocols. Over the years, network scanning has become progressively faster—from 4 months [40] down to 30 days [18], then 1 day [24], and now 45 min [15]. In conjunction with these studies, low-overhead OS fingerprinting can allow significantly better understanding of the systems researchers interact with and improve our general knowledge about the Internet.

OS fingerprinting consists of two approaches—*passive* and *active*. The former [22], [57] monitors ongoing communication (inbound and/or outbound) with remote hosts, but does not generate traffic of its own. Unless each studied device voluntarily connects to the measurement server, this technique is difficult to use for classifying each IP on the Internet. The latter approach, which is the topic of this paper, actively sends packets to targets and infers their operating system from the collected responses.

One important aspect that differentiates between the active methods is the potential maliciousness of probing traffic, where certain nonsensical combinations of TCP flags (e.g., SYN-FIN-RST-ACK) or intrusive actions (e.g., trying to delete the root directory in HTTP fingerprinting [45]) may harm or crash the target. Additionally, these packets are easily detected and dropped by IDS [47], which leads to complaints against research institutions using these methods and possibly reduced accuracy of the results.

The second aspect is the amount of outbound traffic required by the classifier, which ranges from a single SYN probe [4], [53] to lengthy multipacket exchanges [27], [31], [45], [50], [55]. Ideally, fingerprinting should be performed with no extra overhead to scan traffic, which rules out techniques [31], [55] that expect to reach the target on multiple open ports, using different protocols (e.g., ICMP, TCP, UDP), and elicit responses on closed ports. While LAN environments can tolerate high traffic rates and may allow multiprotocol access to each host, these conditions are generally difficult to satisfy when scanning the entire Internet.

The third aspect is the ability of the underlying estimator to correctly identify the target OS under realistic network conditions and without using retransmission. Since prior single-packet techniques [4], [53] were mainly developed for local use, they are not well provisioned to overcome high amounts of fluctuation and loss in temporal features. They also lack resilience to OS tuning, which can be applied by end-users in hopes of optimizing network performance or obfuscating the default parameters of the stack. Either way, the modified OS features may

1063-6692 © 2015 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

exhibit little correlation to those originally present at the host, which cripples estimation accuracy of existing tools.

### B. *Contributions and Implications*

Given the many open issues in wide-scale fingerprinting and lacking performance analysis in the literature, our first goal is to formalize the estimation problem in single-packet OS classification and study the pitfalls of existing techniques. We then develop a low-overhead framework we call *Hershel*[1] for overcoming the various randomization effects (i.e., queuing delays, packet loss, manual tuning) and apply it as proof-of-concept in a measurement study that classifies every visible Web server on the Internet.

We next discuss the ethical implications of this work. Our main objective is to benefit researchers studying the Internet at wide scale and provide a solution to an interesting mathematical problem. However, one may become concerned that intruders can use our algorithms for detection of vulnerable operating systems and better tailor the attack payload to particular configurations (e.g., patch levels) of the targets. As opposed to nmap, our techniques require no additional bandwidth during port scanning, which makes them completely stealthy against IDS and other security monitors.

While hypothetically this may be true, we do not believe there is great cause for concern. With modern botnets, large-scale port scanning can be performed in a highly decentralized fashion, with very little traffic originating from each hijacked IP. This affords the attackers a luxury of using more verbose OS fingerprinting tools (i.e., nmap) and still remaining undetected. Researchers, on the other hand, are typically constrained to a single subnet whose generation of disruptive volumes of highly anomalous traffic is bound to attract negative attention.

Additionally, we are not aware of any evidence confirming that attackers are interested in profiling discovered devices using only SYN packets. Recent studies [58] show that once an open port is found, bots either perform more extensive testing of the open service or attempt all known exploits (some outdated by decades) against the port without discrimination. Eliminating nmap from the picture and directly interacting with the service is much quicker and more informative in that context. We therefore do not see OS fingerprinting as a practical technique for increasing maliciousness of the Internet ecosystem.

## II. Related Work

OS fingerprinting has roots in *banner grabbing*, which relies on application-layer protocols (e.g., HTTP, SSH, SMTP, FTP, telnet) to provide a textual description of the OS as part of the communication sequence. While this worked well 20 years ago, banner grabbing today faces many impediments, including high overhead, administrator ban on OS-identifying strings in responses, generic software (e.g., Apache) that can run on multiple platforms without exposing the underlying OS, and purposefully incorrect banners that aim to mislead the various fingerprinting tools.

[1]William J. Herschel invented forensic usage of fingerprints in 1858.

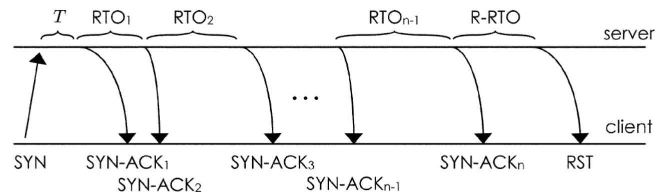


Fig. 1. Retransmission timeouts (RTOs) between SYN-ACK packets.

### A. *Multiple Packets*

The second wave of OS classification started in 1997 with the release of nmap [31], which pioneered TCP/IP tricks that would elicit different responses from different implementations. By default, it sends 1032 probes to the target, including a vertical port scan and certain malformed packets that trigger popular IDS such as Snort [42]. Nmap ideally expects the target to accept a TCP connection, send ICMP port unreachable on a closed UDP port, and respond to a ping. Under bandwidth-optimized settings for OS classification, nmap requires no fewer than 38 different probes. However, due to mandatory retransmission, this in practice corresponds to well over 100 packets per host.

Due to its popularity, nmap has received a great deal of attention in the literature, which includes usage of neural networks to differentiate between versions of the same OS [44], detection of unknown devices [26], and techniques for reducing the number of sent probes [16]. Additional work includes fuzzy matching [55], automatic generation of OS features that aid fingerprinting [9], [41], application of formal testing methods to the detection problem [17], and classification using lengthy observations (up to 100K packets) of initial sequence numbers (ISNs) from the TCP header [27].

Besides the amount of traffic generated by multipacket tools in large-scale scans, another problem is the prevalence of load balancers in the Internet today. These devices, commonly placed in front of servers, may disperse consecutive probes to different physical machines or perform certain elements of the handshake themselves, leading to jumbled fingerprints. This can be avoided by scanning techniques that rely on one outgoing packet, which we describe next.

### B. *Single Packet*

RING [53] and Snacktime [4] are the only tools that perform classification using a single outbound probe. As shown in Fig. 1, each measurement consists of a SYN packet, server processing delay $T$ needed to accept the connection, and a stream of $n$ SYN-ACK responses from the target OS, followed by an optional TCP reset (RST) with its own RTO. RING uses the $n-1$ values in the RTO vector and presence of the final RST packet in classification. Snacktime ignores the RST feature, but instead uses the default TCP window size and TTL carried in the SYN-ACKs, which allows it to differentiate between 25 operating systems [4]. We analyze its classification process in more detail later in the paper. A simplified version of Snacktime and extension to 98 signatures was offered in [20] and [24]. However, no accuracy analysis, modeling, or verified improvement was provided.

Another tool with a related capability is p0f [57]. In addition to passive fingerprinting, it can actively generate SYN packets and profile remote network stacks based on a set of fixed features from the SYN-ACKs (i.e., window size, TTL, IP flags, and TCP options). However, it does not leverage the RTOs and by default is quite verbose (i.e., sends eight copies of the same SYN per target). The current version can differentiate between 18 operating systems.

### C. Common Defenses

There exist many approaches to thwart remote OS fingerprinting. The most basic tools tweak Windows registry [11], [32] or implement plugins [5], [6], [43] for the Unix packet-mangling module Netfilter [29]. Their objective is to modify the fixed features of departing packets to no longer resemble those of the underlying host. A similar direction is to deploy network honeypots [38], [52] or standalone systems [54] that spoof arbitrary operating systems and their services. Placing obfuscation into the network gives rise to intermediate devices known as *fingerprint scrubbers* [37], [46].

While these techniques can effectively deal with static header fields, they are not well suited for distorting the temporal features of departing packets, which requires expensive buffering of packets and per-flow state. Additionally, lack of technical support and possibility for various side-effects (e.g., disabling SACK in TCP may lead to significantly lower throughput) raise questions about deployment of these tools in production systems and/or at large scale.

## III. Stochastic Model

We assume a single-packet scanner similar to Snacktime in Fig. 1. While this approach has minimal intrusiveness, lowest transmission overhead, and nonmalicious operation, it also exhibits several fundamental challenges. These arise due to the complex ways in which the RTOs can be modified by packet traversal across wide-area networks, scarcity of information about the target host contained in the samples, and user tuning of features, all of which have a strong influence on one's ability to detect the underlying OS.

Our contribution in this section is to formalize single-packet OS fingerprinting, set forth clear goals for the classifier, study the impact of network delay and loss on the measured samples, analyze the existing methods, and outline the assumptions under which the classification problem is tractable.

### A. Objectives

Assume a database $\mathcal{D} = (1, 2, \ldots, M)$ of $M \geq 1$ known operating systems, where each OS $j$ has some vector-valued fingerprint $y_j$ collected during *a priori* measurement of the OS. The fingerprint consists of multiple features, which we partition into those modified only by the network (e.g., RTOs) and those only by the user (e.g., TCP window size). Suppose the former are described by some vector $\delta_j$ and the latter by another vector $u_j$. While the length of $\delta_j$ normally depends on $j$, that of $u_j$ is constant across all operating systems.

As both vectors undergo random modification before being observed by the scanner, the response of OS $j$ to probe traffic is some random variable that is a function of $y_j$. Given an observation $x = (\delta; u)$ from an Internet host, a typical estimation problem is to find the most likely OS $s(x)$ that could have produced that vector

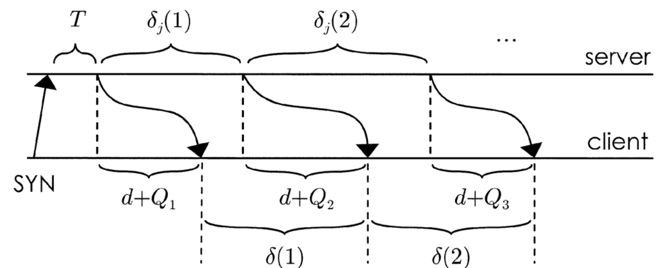

Fig. 2. Effect of jitter on observed RTOs.

$$s(x) := \arg\max_{j\in\mathcal{D}} \; p(y_j|x) = \arg\max_{j\in\mathcal{D}} \frac{p(x|y_j)p(y_j)}{p(x)} = \arg\max_{j\in\mathcal{D}} \; p(x|y_j)p(y_j) \tag{1}$$

where notation $p(x|y)$ refers to the probability (or conditional density, if more convenient) of $x$ given $y$. Observe that the probability $p(x)$ that some OS in $\mathcal{D}$ has produced $x$ is constant for a given observation and can be omitted from the optimization. If the fraction of Internet hosts $p(y_j)$ running OS $j$ is unknown, it is common to set each value to $1/M$, which removes this term from the optimization as well.

The more interesting component of (1) is the probability $p(x|y_j)$ that OS $j$ has produced the observation, or equivalently that $y_j$ has become distorted into $x$. Before investigating this metric further, observe that network and user modifications to the OS features can be treated as independent, from which it follows that

$$p(x|y_j) = p(\delta|\delta_j)p(u|u_j). \tag{2}$$

This means that the two terms can be dealt with separately, which we do in the rest of the section.

### B. Network Features: Jitter

For single-packet techniques [4], [53] in Fig. 1, the vector of temporal features $\delta_j$ consists of individual RTOs generated by network stack $j$. Classification based on $\delta_j$ is possible not only because some devices deviate from TCP algorithms (e.g., exponential timer backoff), but also because RFCs that govern TCP retransmission [8], [35], [36] do not specify the initial RTO or how many SYN-ACKs must be generated. As a result, a wide variety of unique RTO patterns exists.

For the time being, assume loss-free conditions. During collection of sample $x$, suppose $d$ is the sum of propagation and transmission delays along the path from the server back to the scanner. Note that $d$ is a constant due to the fixed size of SYN-ACKs. Now define $Q_m$ to be a random queuing delay of the $m$th packet in the return path. As shown in Fig. 2, the RTO vector $\delta_j$ undergoes distortion that is independent of the forward path, server think time $T$, and propagation delay $d$

$$\delta(m) = \delta_j(m) + Q_{m+1} - Q_m, \; m = 1, 2, \ldots, |\delta_j|. \tag{3}$$

Defining one-way delay (OWD) jitter $J_m = Q_{m+1} - Q_m$ and considering that the gap between subsequent SYN-ACKs

TABLE I
SNACKTIME EXAMPLE

| | $\text{RTO}_1$ (sec) | $Y_{j1}$ | $\text{RTO}_2$ (sec) | $Y_{j2}$ | $W_j$ |
|---|---|---|---|---|---|
| Observation $\delta$ | 3.0 | | 24.0 | | |
| Fingerprint $\delta_1$ | 3.0 | 6 | 12.0 | 0 | 6 |
| Fingerprint $\delta_2$ | 2.9 | 1 | 23.9 | 1 | 2 |

is sufficiently large (i.e., at least several seconds), it follows that back-to-back packets arriving from the server are not likely to encounter the same busy period of the queues they traverse. In that case, it is reasonable to assume that sequence $Q_1, Q_2, \ldots$ consists of independent and identically distributed (i.i.d.) random variables. Furthermore, since the number of hops and congestion of the path is not affected by $j$, the distribution of each $Q_m$ does not depend on the OS being profiled. This leads to

$$p(\delta|\delta_j) = \begin{cases} \prod_{m=1}^{|\delta|} f(\delta(m) - \delta_j(m)), & |\delta| = |\delta_j| \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

where $f(.)$ is the probability density function (PDF) or probability mass function (PMF) of OWD jitter, depending on whether $J_m$ is treated as continuous or discrete. It should also be noted that $Var[J_m] = 2Var[Q_m]$, while $f(.)$ is zero-mean and symmetric. For certain models of OWD, jitter can be obtained in closed form. For example, exponential $Q_m$ produces the Laplace distribution with the same parameter $\lambda$ and Gaussian $N(\mu, \sigma^2)$ becomes $N(0, 2\sigma^2)$.

We next contrast (4) with the RTO classifier in Snacktime [4], which is a tool that is the closest to our objectives and most advanced in single-packet OS fingerprinting. For each RTO $m$, this method first computes the number of matching digits (limited to six decimal places of precision) between the sample and all known fingerprints $j$

$$Y_{jm} = \max(\lceil -\log_{10}(\max(|\delta(m) - \delta_j(m)|, 10^{-6}))\rceil, 0).$$

It then assigns score $W_j$ to OS $j$ using the sum of these weights across all available RTOs

$$W_j = \sum_{m=1}^{|\delta|} Y_{jm}. \tag{5}$$

For the example in Table I, which exemplifies the common pitfalls of Snacktime, (5) scores six for the first OS and two for the second OS, indicating that jitter combination (0, 12) is more likely than (0.1, 0.1). Taking the log of (4), our model can also be reduced to optimizing a summation

$$\log p(\delta|\delta_j) = \sum_{m=1}^{|\delta|} \log f(\delta(m) - \delta_j(m)). \tag{6}$$

However, it differs from (5) in two important ways. First, the log is applied to the distribution function $f(.)$ rather than the jitter itself. Second, there is no loss of precision due to rounding to the nearest integer or capping the jitter at $10^{-6}$.

Nevertheless, while (4) is a good starting point, it does not work in real networks due to the lacking robustness against packet loss. This is our next topic.

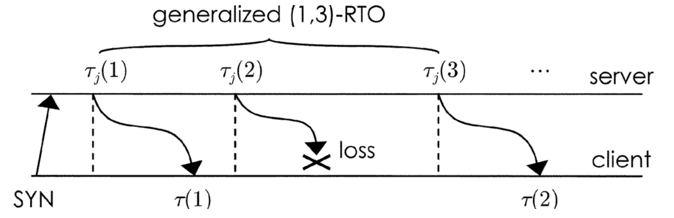


Fig. 3. Generalized RTOs under packet loss.

### C. Network Features: Loss

The main problem with (4) is that loss-free conditions are impossible to satisfy during Internet scans. Besides congestion, routing loops, and various checksum violations, the RTOs may be altered by the target server crashing or shutting down during the measurement, which affects the tail of the RTO vector and appears similar to packet loss. Since single-packet fingerprinting *by definition* cannot retransmit SYN probes, OS detection must be performed using only the features available in observation $x$, which calls for more sophistication in the model.

To exacerbate the situation, packet loss creates more dramatic changes to the RTO vector than delay jitter. For example, consider a scenario with $\delta_j = (3, 6, 12)$, where all delays are given in seconds. Even with a relatively large $E[Q_m] = 100$ ms, delay jitter remains small compared to each RTO. On the other hand, the loss of a single packet produces one of four dissimilar combinations—(3, 6), (3, 18), (6, 12), or (9, 12)—while that of two packets leads to six additional options—(3), (6), (9), (12), (18), or (21). The RTO swing in these cases is significantly higher, which makes mapping $x$ to the correct OS more challenging.

We now examine how to model the combined probability that loss and jitter transform $\delta_j$ into observation $\delta$. This will allow us to solve such dilemmas as whether $\delta = (3, 18)$ is a more likely match to (3, 6, 12) with one lost packet or to some other signature (2.6, 17.9) without any loss. To deal with these cases, we propose to generalize the concept of RTO. First, let $\tau_j$ be a vector of $|\delta_j| + 1$ packet-transmission timestamps from OS $j$

$$\tau_j(m) = \begin{cases} 0, & m = 1 \\ \tau_j(m-1) + \delta_j(m-1), & m \geq 2 \end{cases} \tag{7}$$

and $\tau$ be the corresponding random vector observed in $x$ after the packets have traversed the network. Then, a generalized $(m, m+k)$-RTO is the distance $\tau_j(m+k) - \tau_j(m)$, which is illustrated in Fig. 3 for $m = 1$ and $k = 2$. Note that $k = 1$ produces the usual RTO and that all timestamps are given using local clocks (i.e., $\tau_j$ at the server and $\tau$ at the client).

Now suppose set $\Gamma(\tau, \tau_j)$ contains all subsets of size $|\tau|$ of integer sequence $(1, 2, \ldots, |\tau_j|)$. We can view each $\gamma \in \Gamma(\tau, \tau_j)$ as a mapping of received packets in $\tau$ to their position in the original vector $\tau_j$, i.e., $\gamma(m) = k$ means that the $m$th received SYN-ACK was initially in position $k$. For the example in Fig. 3, we have $\gamma(1) = 1$ and $\gamma(2) = 3$. Assuming no reordering of SYN-ACKs, which is reasonable given at least several seconds between them, each $\gamma$ is a vector of strictly increasing integers.

Armed with these definitions, we get

$$p(\tau|\tau_j) = \begin{cases} \sum_{\gamma \in \Gamma(\tau, \tau_j)} p(\gamma) p(\tau|\tau_j, \gamma), & |\tau| \leq |\tau_j| \\ 0, & \text{otherwise} \end{cases} \tag{8}$$

where the number of summation terms equals the number of ways to select $|\tau|$ objects from $|\tau_j|$ available options and (8) is nonzero only if the number of received packets does not exceed that in the fingerprint. This is in contrast to (4), where the two vectors had to have equal length.

Again leveraging the large spacing between server responses, we can treat congestion events affecting SYN-ACKs as independent, which allows one to approximate packet loss as an i.i.d. Bernoulli process with some probability $q$. Since each loss combination is equally likely, we get

$$p(\gamma) = q^{|\tau_j|-|\tau|}(1-q)^{|\tau|} \quad (9)$$

which can be moved outside the summation in (8). To deal with $p(\tau|\tau_j, \gamma)$, which is the probability to observe $\tau$ from OS $j$ under loss pattern $\gamma$, notice that the gap between each adjacent pair of received packets is determined by the generalized RTO

$$\tau(m) - \tau(m-1) = \tau_j(\gamma(m)) - \tau_j(\gamma(m-1)) + J'_m \quad (10)$$

where $m \geq 2$ and generalized jitter $J'_m$ is given by

$$J'_m = Q_{\gamma(m)} - Q_{\gamma(m-1)}. \quad (11)$$

Rearranging the terms in (10), define the $m$th jitter sample under pattern $\gamma$ as

$$R^{\gamma}_{jm} = \tau(m) - \tau(m-1) - \tau_j(\gamma(m)) + \tau_j(\gamma(m-1)). \quad (12)$$

Noticing that $J'_m$ has the same distribution as $J_m$ yields

$$p(\tau|\tau_j, \gamma) = \prod_{m=2}^{|\tau|} f(R^{\gamma}_{jm}). \quad (13)$$

We thus get for $|\tau| \leq |\tau_j|$

$$p(\tau|\tau_j) = q^{|\tau_j|-|\tau|}(1-q)^{|\tau|} \sum_{\gamma \in \Gamma(\tau,\tau_j)} \prod_{m=2}^{|\tau|} f(R^{\gamma}_{jm}) \quad (14)$$

which replaces $p(\delta|\delta_j)$ in (2).

### *D. User Features*

OS tuning is common practice in the current Internet, with numerous online guides recommending optimizations to network settings [33], [51] and automated software offering tuning capabilities to the TCP/IP stack to achieve better performance [14]. A number of fixed header parameters in general-purpose kernels (e.g., Unix, Windows) can be changed through registry or using command-line tools. Unlike jitter-induced noise, where small distortions are generally more likely than large ones, the main difference with OS tuning is that *there may be no correlation between the manually selected values of the user and those installed in the OS by default*. For example, TCP window size may be more likely to jump from 8192 to 65 535 than to 8193.

While accurate modeling of manual modification and human psychology is difficult, it makes sense for the analysis to at least take into account whether a given feature under user control has been changed. Suppose that $\pi_m$ is the probability of such modification in feature $m$ and assume that user tuning is applied independently to each available parameter. Defining $I_{jm} = \mathbf{1}_{\{u(m)=u_j(m)\}}$ to be an indicator of the event that the $m$th measured feature matches the original of OS $j$, we get

$$p(u|u_j) = \prod_{m=1}^{|u|} \Big[(1-\pi_m)I_{jm} + \pi_m(1-I_{jm})\Big]. \quad (15)$$

Besides user interference, vector $u_j$ may be modified by intermediate devices along the path (e.g., NAT, IDS, fingerprint scrubbers [11], [37], [43], [46], [54]), whose actions can be clumped under the same umbrella of (15). Since buffering packets for periods of time comparable to RTO (i.e., 3–6 s) and per-flow state are expensive, it is often safe to assume that these devices do not alter the RTO pattern in significant ways and thus leave enough features by which the OS can still be identified. This underscores the importance of having a robust RTO estimator.

The Snacktime algorithm for scoring user-modified features can be generalized as a sum of weights assigned to each match

$$W'_j = \sum_{m=1}^{|u|} w_m I_{jm} = \sum_{I_{jm}=1} w_m \quad (16)$$

which is added to the RTO score $W_j$ in (5) for a final result. One open issue, however, is selection of proper weights, which need to be somehow correlated with feature volatility. Our model is much simpler since $\pi_m$ directly provides this probability. To better understand the difference between (15) and (16), assume that $\pi_m > 0$ for all $m$ and write

$$\log p(u|u_j) = \sum_{I_{jm}=1} \log(1-\pi_m) + \sum_{I_{jm}=0} \log \pi_m. \quad (17)$$

For $\pi_m \approx 1$, we get $\log \pi_m \approx 0$, the second term of (17) disappears, and our model reduces to Snacktime with weights $w_m = \log(1-\pi_m)$. However, in more realistic cases of $\pi_m \ll 1$, the second term of (17) becomes nonnegligible and serves the role of balancing nonmatching features against those that do match. Snacktime has no such mechanism.

### *E. Final Result*

We now consolidate the various models into one formula. Combining (14) and (15) in (2) and (1), dropping terms that do not depend on $j$, and performing straightforward manipulations, we get

$$s(x) = \underset{j \in \mathcal{D}:|\tau| \leq |\tau_j|}{\arg\max} \Bigg\{ p(y_j) q^{|\tau_j|-|\tau|} \sum_{\gamma \in \Gamma(\tau,\tau_j)} \prod_{m=2}^{|\tau|} f(R^{\gamma}_{jm}) \times \prod_{I_{jm}=1} (1-\pi_m) \prod_{I_{jm}=0} \pi_m \Bigg\}. \quad (18)$$

Although (18) maximizes the OS-detection probability under the assumptions stated throughout this section, its performance with *a priori* unknown $q$, $\pi_m$, $f(.)$, and $p(y_j)$ is an open question. We return to it later in the paper; in the meantime, we outline the various remaining issues.

### *F. Limitations*

First, the SYN packet may be lost and never reach the target. Since there is no way to verify this, the host will automatically be considered nonresponsive and will be excluded from fingerprinting. Not much can be done to overcome this problem unless SYN retransmission is allowed. If we relax the single-packet assumption, the estimator will face the problem of determining which of the SYNs triggered which SYN-ACK

TABLE II
SAMPLE SIGNATURES (M = MSS, N = NOP, W = Window Scale, S = SACK, T = Timestamp)

| Operating system | Win | TTL | DF | Reset | | | | MSS | OPT | SA-RTO | R-RTO |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | RST | RA | RN | RW | | | | |
| Windows 7 | 8192 | 128 | 1 | 1 | 0 | 1 | 0 | 1460 | MNWST | 3, 6 | 12 |
| Linux 2.6 | 5792 | 64 | 1 | 0 | – | – | – | 1460 | MSTNW | 3.8, 5.9, 12.1, 24, 48.2 | – |
| Linux 2.0 | 32736 | 64 | 0 | 0 | – | – | – | 1414 | M | 3, 6, 12, 24, 48, 96 | – |
| Mac OS 10.3 | 33304 | 64 | 1 | 1 | 1 | 1 | 32768 | 1460 | MNWNNT | 2.92, 6, 12, 24 | 30 |
| NetBSD 4.0.1 | 32768 | 64 | 1 | 0 | – | – | – | 1460 | MNWNNTSNN | 2.92, 6, 12, 24 | – |
| VxWorks 5.4.2 | 8192 | 64 | 0 | 1 | 1 | 1 | 8192 | 512 | MNW | 5.58, 24 | 45 |
| Juniper Netscreen | 8192 | 64 | 0 | 1 | 0 | 0 | 8192 | 1380 | M | 1.67, 2, 2, 2, 2, 2, 2, 2 | 2 |

TABLE III
EXAMPLES OF TRANSFORMATIONS APPLIED BY THE OS TO TCP OPTIONS (DASHES INDICATE IMPOSSIBLE CASES)

| Operating system | All enabled | Drop S | Drop T | Drop W | Drop ST | Drop SW | Drop WT | Drop all |
|---|---|---|---|---|---|---|---|---|
| Linux 2.6 | MSTNW | MNNTNW | MNNSNW | MST | MNW | MNNT | MNNS | M |
| Windows XP/2003 | MNWNNTNNS | MNWNNT | MNWNNS | MNNTNNS | MNW | MNNT | MNNS | MNW |
| Windows 7/2008 | MNWST | – | – | MST | – | – | – | – |
| FreeBSD 8.2 | MNWST | MNWNNT | – | – | – | – | MSE | M |
| Solaris 10 | NNTMNWNNS | NNTMNW | – | – | – | – | – | – |

response, without which the RTOs cannot be computed correctly. This problem can be solved in the future by encoding the retransmission attempt into the source port of the SYN.

Second, our model allows only the *network* to modify the received RTOs. However, this may not hold if users manage to alter SYN-ACK spacing during OS tuning. This is not of widespread concern as few optimization guides target the RTO pattern. With enough effort, scrubbers and obfuscation tools can disrupt inter-SYN-ACK delays. However, we do not consider development of end-to-end methods to combat such approaches a fruitful objective. A related problem arises with middleboxes and caches that accept the connection on behalf of the server [19], in which case any fingerprinting tool is bound to classify only the visible side of the TCP stream (i.e., the OS of the middlebox).

Third, Hershel's accuracy may deteriorate if the network jitter process $J_m$ becomes non-i.i.d. or deviates from the predicted bounds, e.g., due to significant kernel scheduling latency during CPU overload. Similar issues may surface if network loss depends on $j$, users modify different operating systems with different probability, or there is correlation in loss events within a single stream of SYN-ACKs. Solving these problems requires a per-OS set of parameters $(q_j, f_j(.), \pi_{jm})$ and multidimensional covariance matrices for joint distributions of RTOs and individual features modified by tuning, all sampled under realistic load conditions. Needless to say, these are difficult to come by, but we may consider this direction in future work.

## IV. CLASSIFIER

Our next contribution is to enhance Snacktime's feature vector, describe a working classifier based on the theory developed in Section III, bring awareness to RTO randomization performed by certain OSs, and explain how to collect signature databases under these conditions.

### A. Features

Snacktime uses only two non-RTO features—TCP advertised window size and TTL. However, additional parameters are readily available from the SYN-ACKs. Following Table II, these include the Do Not Fragment (DF) flag in the IP header, four different fields from the RST packet (more on this below), the Maximum Segment Size (MSS) declared by TCP, the order in which the OS assembles the option fields (OPT), SYN-ACK RTOs (SA-RTO), and the RST RTO (R-RTO). Some of these features are self-explanatory, but others require additional elaboration.

First, it should be noted that the initial TTL cannot be reconstructed exactly at the receiver. We use the common technique of rounding this value up to the nearest "likely" boundary, which includes four values used by the OSs in our database $\mathcal{D}$—32, 64, 128, and 255. Second, the reset features are quite rich. In Table II, the binary flag RST is 1 for the fingerprints that contain a reset packet, RA indicates whether the RST has the ACK bit set, RN is 1 if the ACK sequence is nonzero, and RW records the window of the reset packet. RST features represent peculiarities of internal stack operation and cannot be modified via OS tuning. However, fingerprint scrubbers, NAT/IDS, and kernel recompilation can still change them.

Third, as seen in the table, support for TCP options differs between the operating systems since no specific subset is required to be implemented [21]. More importantly, users have the freedom to disable them as needed. As certain options are considered security risks (e.g., timestamps), they may be disabled by default, although users can still reenable them. Certain devices (e.g., printers) do not allow OPT tweaking at all, while newer versions of popular operating systems tend to support fewer choices. For example, even though Windows 7/2008 provides registry keys to disable TCP timestamps, the modification does not work. Similarly, SACK can be disabled only if the entire TCP stack is offloaded to the NIC [28].

What makes OPT a good feature is not the specific string, but rather *the order in which non-padding options appear*. This is illustrated in Table III, where we progressively disable various combinations of options and observe the resulting SYN-ACK packets. For example, Windows XP supports four options MWTS. Turning off W produces MTS interspersed by NOPs as padding. Simplicity of implementation and lacking reasons to reorder the options suggests that this phenomenon likely exists in other stacks.

TABLE IV
ENHANCED FEATURE VECTOR

| Feature | Description | Appeared In |
|---|---|---|
| Win | Receiver window | [4], [31], [50], [55], [57] |
| TTL | Time-to-live field | [4], [31], [50], [55], [57] |
| DF | Do Not Fragment | [31] [50] [55], [57] |
| SA-RTO | RTO sequence | [4], [50], [53] |
| RST | True if RST packet | [53] |
| MSS | Max segment size | [31], [50], [57] |
| OPT | TCP options (exact) | [31], [57] |
| RA | ACK bit in RST | New |
| RN | ACK seq $\neq 0$ in RST | New |
| RW | Window in RST | New |
| OPT | TCP options (subset) | New |
| R-RTO | RTO of RST packet | New |

As a result, OPT requires a more advanced classification logic than straight comparison. Specifically, a match is registered if the observed sample $x$ contains a *feasible* string, which we examine by taking an intersection of non-NOP options between $x$ and each fingerprint, followed by verification that the order of the resulting letters is the same. For example, MTW is a match to Linux, VxWorks, and Juniper in Table II, but not the other OSs.

Fourth, the reset RTO (R-RTO) helps in resolving additional ambiguities, such as between Mac OS 10.3 and NetBSD 4.0.1 in Table II, which otherwise have identical SA-RTO patterns. Additionally, we expand Snacktime's default measurement time limit from 65 to 120 s, the latter of which is the maximum segment lifetime (MSL) of TCP [36]. For instance without considering the 96-s RTO of Linux 2.0 in Table II, it might be hard to differentiate it from Linux 2.6.

Table IV summarizes the features used in our classification and compares them to those in nmap, p0f, and Xprobe [31], [50], [53], [55], [57]. We have four novel features and one match type (subset) never used in fingerprinting before.

## B. Stochastic Timers

Table II shows SA-RTOs from a single captured sample of the OS. However, it turns out that many kernels naturally exhibit significant RTO variation, sometimes by as much as 50%. Two examples are shown in Fig. 4 using a 2-D scatter plot of the first two SA-RTOs. For Server 2003 in Fig. 4(a), there are two distinct patterns—the lower left corner, with $\text{RTO}_1$ distributed in $[2.2, 3.3]$ and $\text{RTO}_2$ frozen at 6.56, and the upper section, with $\text{RTO}_1$ scattered in $[3.3, 4.6]$ and $\text{RTO}_2$ in $[9.5, 9.8]$. Furthermore, the two scenarios are not equally likely as the bottom one occurs 68% of the time. This shows that the temporal model must take into account not just the possible RTO regions, but also their likelihoods.

A similar picture emerges for Linux 2.6 in Fig. 4(b). The mass of the RTO is now concentrated on 11 distinct points, where $\text{RTO}_1$ ranges from 3 to 4.4 s and $\text{RTO}_2$ from 6 to 6.2. Again, the popularity of individual points is nonuniform, swinging from 2% to 16%. Note that both cases in Fig. 4 have been collected from idle hosts over a single-hop network consisting of one switch, *which makes this behavior part of the fingerprint itself rather than an artifact of the sampling environment*.

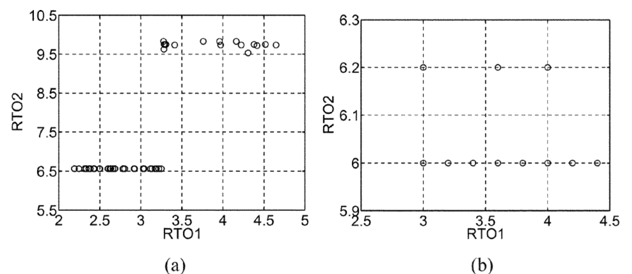


Fig. 4. RTO randomness in TCP/IP scheduler. (a) Windows Server 2003. (b) Linux 2.6.

Possible reasons for this fluctuation are the absence of per-connection RTO timers during the SYN-ACK phase and discretization of retransmission delays. What these examples show is that internal OS operation is a complex stochastic system that requires measuring the RTO *distribution* (rather than a single snapshot) during creation of the signature database. This is necessary because such large variations are not taken into account by the jitter model, which normally assumes OWDs on the order of tens or hundreds of milliseconds, with similarly sized jitter.

Our approach is to treat RTOs as random variables, unlike prior work that has always considered them deterministic. Specifically, suppose OS $j$ has $w_j$ unique types of behavior, each occurring with probability $\beta_{jr}$, where $r = 1, 2, \ldots, w_j$. We call each of these types a *subOS* and assign it a separate RTO vector $\tau_{jr}$, which updates (13) to

$$p(\tau|\tau_j, \gamma) = \sum_{r=1}^{w_j} \beta_{jr} p(\tau|\tau_{jr}, \gamma). \tag{19}$$

A simpler technique is to measure each host $w$ times and let each obtained RTO vector $\tau_{jr}$ be a subOS with $\beta_{jr} = 1/w$. In that case, (19) becomes

$$p(\tau|\tau_j, \gamma) = \frac{1}{w} \sum_{r=1}^{w} \prod_{m=2}^{|\tau|} f(R_{jrm}^{\gamma}) \tag{20}$$

where $R_{jrm}^{\gamma}$ is the generalized jitter of the $m$th RTO under subOS $r$ of OS $j$ and loss pattern $\gamma$. Note that summations involving $\Gamma(\tau, \tau_j)$ remain the same since all subOSs within a given OS send a fixed number of SYN-ACKs. They also exhibit deterministic user features, which keeps (15) unchanged.

## C. Fingerprint Database

In order to produce an accurate fingerprint $\tau_j$, the OS must be measured in some isolated testbed with low end-to-end delays and idle conditions at the server. To avoid loss-related bias, each host must be sampled multiple times to determine the longest vector of RTOs it produces, which should then be used to collect $w$ loss-free samples for the database. Following these guidelines, we installed a variety of commodity operating systems in our lab, determined the proper size of their RTO vectors, and collected $w = 50$ clean fingerprints from each. We also captured a number of embedded devices found in our department LAN.

TABLE V
CLASSIFICATION ACCURACY (PERCENT) OF ISOLATED FEATURES

| Win | TTL | DF | OPT | MSS | RST |
|---|---|---|---|---|---|
| 39.2 | 4.7 | 1.8 | 1.1 | 5.5 | 14.4 |

The final step was to perform a scan of the university network for additional signatures not already in the database. Once found, these devices were fingerprinted in a similar fashion, producing $w$ clean samples per OS. When the owner of the device could not be contacted and the server's Web page did not provide enough detail, we used nmap to identify the device. While Snacktime ships with 25 signatures and [24] uses 98, our database contains 116 network stacks. We can distinguish not only between different operating systems (e.g., Windows, Linux, FreeBSD), but also sometimes identify their versions and patches (e.g., Windows Server 2003 with and without SP1, MacOS 10.3 versus MacOS 10.4).

### D. Hershel

Our classification method, which we call *Hershel*, builds upon (18) and (20), where we treat all $w = 50$ subOSs as deterministic. Common sense suggests that users, scrubbers, and network devices are not likely to directly tweak individual RST features RA, RN, and RW; instead, these fields (if modified at all) will be simultaneously replaced with another set that comes from a different OS. We thus combine all four RST values in Table II into one atomic feature for classification purposes. This makes vector $u_j$ consist of six fields—Win, TTL, DF, MSS, OPT, and aggregated RST. Table V shows the accuracy of individual features across the entire database (all ties are broken uniformly randomly).

RTO vectors $\tau$ and $\tau_j$ include timestamps of all SYN-ACKs and the first RST (if present). To account for resets that might be injected by firewalls/IDS after they time out the connection, (8) and (15) require a revision. Specifically, if the measured vector $\tau$ contains a reset, but $\tau_j$ does not, the RST is removed from $\tau$ prior to computing (8). To account for the mismatch in the RST feature, (15) gets multiplied by $\pi_6$. In the opposite case, i.e., $\tau_j$ contains a RST, but $\tau$ does not, it is important to avoid mistaking packet loss for changes in the RST feature and improperly penalizing $p(u|u_j)$ with $\pi_6$. Next, if both vectors contain a reset packet, (15) gets hit with either $\pi_6$ or $1 - \pi_6$ depending on the match in (RA, RN, RW). Finally, if neither vector has a RST, then (15) enjoys multiplication by $1 - \pi_6$.

## V. SIMULATIONS

Our contribution in this section is to explain how to select the parameters of the model and examine Hershel's accuracy in simulations in comparison to Snacktime.

### A. Parameters

For lack of a better assumption, we suppose that all OSs are equally likely to appear in the trace and set $p(y_j) = 1/M$ to be a uniform PMF. While it is possible to consider multiple iterations and refine this value after each pass, the resulting system sometimes exhibits instability and divergence into inferior states. We thus leave stability analysis for future work and only perform a single iteration in our evaluation below.

We use $\pi_m = 0.01$ for RST and OPT, while keeping $\pi_m = 0.1$ for the other features. The rationale is that RST behavior and option ordering can be changed only through kernel source-code modifications and usage of aggressive intermediate devices, neither of which we believe is that common in today's Internet compared to stack tuning. For queuing delay, we use a simple exponential distribution with CDF $1 - e^{-\lambda x}$ whose mean is set to 0.5 s (rate $\lambda = 2$). This produces Laplace jitter density

$$f(z) = \frac{\lambda}{2} e^{-\lambda |z|}. \tag{21}$$

Note that usage of $\lambda = 2$ is fairly pessimistic, with the majority of paths likely exhibiting significantly smaller delays. For example, this model assumes 82% of the paths produce over 100 ms queuing delays, 37% over 500 ms, and 14% over 1 s. For packet loss, we use Google's study [10] to set $q = 3.8\%$, which was their highest rate of SYN-ACK loss.

### B. Results

Our next goal is to examine Hershel's robustness in the presence of OWD jitter, packet loss, and random feature modification by the user. We also aim to assess the sensitivity of results to our choices of default parameters above. We simulate a FIFO queue between the server and the client with a given delay distribution. Each packet is dropped by the router with some probability $q_{\text{real}}$ and each feature is independently modified with another probability $\pi_{\text{real}}$. Since these are per-packet and per-feature metrics, it also makes sense to examine the fraction $\chi = E[(1 - q_{\text{real}})^{|\tau_j|}](1 - \pi_{\text{real}})^6$ of all generated samples that do not have any loss or feature modification, where the expectation is taken over all $j$.

The distribution of popularity $p_{\text{real}}(y_j) \sim j^{-\alpha}$ is set to Zipf with shape parameter $\alpha = 1.2$, which approximates the fact that some OSs are much more popular than others. We do not attempt to make our assignment of index $j$ to each physical OS such that its $p_{\text{real}}(y_j)$ closely follows that in the Internet (which is unknown anyway); instead, the simulation simply verifies performance of the proposed estimator when the OS frequency is highly nonuniform. For that purpose, random ordering of OSs in the database is sufficient.

Table VI shows classification accuracy for several scenarios of interest. We examine three types of OWD with mean $\mu$ in the first column—Pareto $1 - (1 + x/\beta)^{-\alpha}$ with $\alpha = 3$ and $\beta = \mu(\alpha - 1)$, exponential with rate $1/\mu$, and uniform in $[0, 2\,\mu]$. We use the original Snacktime since the simplified version from [20] performs worse. Using just the RTOs, Snacktime in the table starts at close to 13%, but then deteriorates below 1% near the bottom. This amounts to essentially guessing across the 116 available options (i.e., $1/116 = 0.86\%$). Augmented with Win and later TTL, Snacktime begins at a more healthy 52%–58%, but then eventually reduces to single digits.

The next six columns show Hershel with its default $\lambda = 2$. Classifying just based on the RTO vector, Hershel doubles Snacktime's accuracy in the first three scenarios (i.e., the first 12 rows of the table), triples it in the next one, and improves by an order of magnitude in the last one. As additional features are added, Hershel becomes even better, with significant gains

TABLE VI
CLASSIFICATION ACCURACY (PERCENT) IN SIMULATIONS OF $2^{18}$ SAMPLES

| OWD distribution | $\mu$ (sec) | Snacktime RTO | Snacktime +Win | Snacktime +TTL | Hershel $\lambda=2, q=0.038$ RTO | +Win | +TTL | +DF | +OPT | +MSS | +RST | Hershel $\lambda=10$ | Hershel $q=0.1$ |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| $q_{real}=0, \pi_{real}=0$ ($\chi=100\%$) | | | | | | | | | | | | | |
| Pareto | 0.5 | 12.6 | 51.8 | 58.3 | 22.1 | 81.4 | 86.2 | 88.5 | 96.2 | 99.72 | 99.72 | 94.62 | 99.69 |
| Exp | 0.5 | 12.8 | 51.8 | 58.3 | 21.9 | 82.6 | 86.9 | 89.4 | 96.5 | 99.92 | 99.94 | 96.21 | 99.82 |
| Uniform | 0.5 | 13.0 | 51.9 | 58.4 | 21.7 | 84.1 | 87.4 | 89.8 | 96.8 | 99.99 | 99.99 | 98.50 | 99.99 |
| Pareto | 0.1 | 16.3 | 56.9 | 62.9 | 33.1 | 93.0 | 94.9 | 96.7 | 99.0 | 99.99 | 99.99 | 99.69 | 99.99 |
| $q_{real}=3.8\%, \pi_{real}=0$ ($\chi=84\%$) | | | | | | | | | | | | | |
| Pareto | 0.5 | 10.0 | 43.4 | 49.0 | 21.4 | 78.5 | 85.1 | 87.7 | 96.1 | 99.69 | 99.69 | 94.68 | 99.66 |
| Exp | 0.5 | 10.1 | 43.4 | 49.0 | 21.5 | 80.1 | 85.6 | 88.1 | 96.3 | 99.76 | 99.82 | 96.21 | 99.80 |
| Uniform | 0.5 | 10.3 | 43.4 | 49.0 | 21.7 | 81.1 | 86.4 | 89.0 | 96.7 | 99.96 | 99.96 | 98.50 | 99.96 |
| Pareto | 0.1 | 13.1 | 47.9 | 53.2 | 31.6 | 89.6 | 93.6 | 95.6 | 98.8 | 99.96 | 99.96 | 99.66 | 99.97 |
| $q_{real}=3.8\%, \pi_{real}=10\%$ ($\chi=49\%$) | | | | | | | | | | | | | |
| Pareto | 0.5 | 10.0 | 39.9 | 44.4 | 21.4 | 72.7 | 77.7 | 78.6 | 91.4 | 94.93 | 95.37 | 90.13 | 95.25 |
| Exp | 0.5 | 10.1 | 39.9 | 44.4 | 21.5 | 73.8 | 78.3 | 79.1 | 91.6 | 95.02 | 95.55 | 91.78 | 95.34 |
| Uniform | 0.5 | 10.3 | 39.9 | 44.4 | 21.7 | 75.1 | 78.9 | 79.7 | 91.9 | 95.20 | 95.63 | 93.97 | 95.57 |
| Pareto | 0.1 | 13.1 | 44.3 | 48.5 | 31.6 | 83.8 | 87.3 | 87.7 | 95.0 | 96.54 | 96.92 | 96.67 | 96.87 |
| $q_{real}=10\%, \pi_{real}=10\%$ ($\chi=34\%$) | | | | | | | | | | | | | |
| Pareto | 0.5 | 6.9 | 29.9 | 33.4 | 20.1 | 68.1 | 76.2 | 77.1 | 91.2 | 94.84 | 95.22 | 90.01 | 95.14 |
| Exp | 0.5 | 7.0 | 29.9 | 33.4 | 20.1 | 69.2 | 76.8 | 77.7 | 91.5 | 94.98 | 95.43 | 91.76 | 95.20 |
| Uniform | 0.5 | 7.2 | 29.9 | 33.4 | 20.1 | 70.4 | 77.4 | 78.3 | 91.7 | 95.13 | 95.51 | 93.82 | 95.46 |
| Pareto | 0.1 | 9.3 | 33.5 | 36.8 | 29.4 | 78.4 | 85.3 | 85.7 | 94.5 | 96.38 | 96.71 | 96.46 | 96.67 |
| $q_{real}=50\%, \pi_{real}=50\%$ ($\chi=0.13\%$) | | | | | | | | | | | | | |
| Pareto | 0.5 | 0.82 | 2.37 | 2.49 | 10.4 | 23.7 | 28.1 | 35.6 | 53.7 | 56.65 | 59.95 | 58.95 | 60.23 |
| Exp | 0.5 | 0.83 | 2.37 | 2.49 | 10.5 | 24.1 | 28.4 | 35.9 | 53.8 | 56.74 | 60.12 | 60.40 | 60.31 |
| Uniform | 0.5 | 0.84 | 2.37 | 2.49 | 10.6 | 24.5 | 28.6 | 36.5 | 54.0 | 56.89 | 60.25 | 60.79 | 60.46 |
| Pareto | 0.1 | 1.11 | 2.90 | 2.95 | 14.4 | 28.3 | 32.0 | 40.5 | 56.8 | 59.45 | 62.68 | 64.84 | 63.06 |

seen at the Win and OPT boundaries. This shows that unlike DF, option strings form an orthogonal dimension to Win/TTL. The MSS improves the result further by 3% and the RST packet by an additional 0.5%–3%, with the impact mostly limited to high-loss cases.

Staying with $\lambda = 2$, observe that Hershel is quite insensitive to selection of $f(z)$. Specifically, classification accuracy improves *not* when $\lambda$ equals $1/\mu$ or the PDF of real delay matches (21), but *as $\mu$ gets smaller or the tail of the delay gets lighter*. This can be seen by contrasting the two Pareto cases ($\mu = 0.1$ and $\mu = 0.5$) and comparing Pareto, exponential, and uniform cases (all with $\mu = 0.5$). As the difference between the last three scenarios is quite small, we conclude that the distribution of network jitter, as opposed to its mean, generally has a minor effect on accuracy. Therefore, keeping the Laplace model (21) appears reasonable.

To shed additional light on selection of parameters, the next column of the table reruns Hershel with all available features and $\lambda = 10$. While this slightly improves the $\mu = 0.1$ case, this happens only under 50% packet loss and at the expense of significant reduction in accuracy in other rows, which suggests that $1/\lambda$ should *overestimate*, rather than *underestimate*, the real network delay. To this end, our previous conservative choice $\lambda = 2$ seems quite appropriate. The last column of the table reverts to $\lambda = 2$ and demonstrates that the model is insensitive to selection of $q$. We thus keep $q = 3.8\%$ for the Internet classification below.

# VI. EXPERIMENTS

Our contribution in this section is to apply Hershel to a wide-scale Internet scan and provide an assessment of the obtained classification.

## A. Dataset Properties

We use Internet scan data from [20], which is based on a 2010 survey of Web servers in [24]. These IPs were discovered by sending port-80 SYN packets from Windows Server 2008 (with all TCP options enabled) to every address in BGP. The experiment garnered 37.8M samples $x$ that contained at least one SYN-ACK, which we later feed into Hershel. We start by examining occurrence of various features in the dataset and their mapping to signatures in $\mathcal{D}$. We qualitatively group them into four types—linux, windows, embedded (routers, modems, cameras, hardware gadgets), and other (BSD, Mac, AIX, NetApp, Big-IP, SunOS).

To first step is to ensure that packet loss has not produced totally unworkable temporal features in the dataset. Table VII shows the number of available RTOs per destination. It is encouraging to see that the top four spots retain enough information for a meaningful match and the most difficult case (i.e., single SYN-ACK) follows in sixth place. While the average number of received packets was 5, one host transmitted over 3M SYN-ACKs. We next analyze sanity of the remaining features and build intuition for what to expect from Hershel classification.

The scan contains a staggering 3815 unique window sizes, while our fingerprint collection $\mathcal{D}$ has only 51. While users tuning their stacks and scrubbers modifying the OS signature are possible reasons, we also found that the advertised window of SYN-ACKs can be easily changed at the application layer by resizing the socket buffer (i.e., calling `setsockopt` with the `SO_RCVBUF` option) before the connection is accepted. This highlights the need for a flexible classifier that allows features to mismatch.

TABLE VII
TOP RTO COUNTS (99% OF TOTAL)

| RTOs | Hosts | Sigs | Group |
|---|---|---|---|
| 3 | 9,639,810 | 27 | all |
| 2 | 9,070,991 | 16 | windows, embedded |
| 5 | 7,834,027 | 23 | linux, embedded, other |
| 4 | 5,066,940 | 16 | unix, embedded |
| 1 | 2,669,222 | 1 | Dell printer |
| 0 | 1,992,196 | 0 | – |
| 6 | 540,042 | 9 | linux, embedded, other |
| 19 | 202,733 | 2 | embedded |
| 18 | 162,442 | 0 | – |
| 17 | 110,335 | 0 | – |

TABLE VIII
TOP WINDOW SIZES (87% OF TOTAL)

| Window | Hosts | Sigs | Group |
|---|---|---|---|
| 5,792 | 10,143,772 | 4 | linux |
| 16,384 | 7,051,858 | 6 | windows, embedded, other |
| 8,192 | 4,266,370 | 17 | windows, embedded |
| 65,535 | 3,551,640 | 9 | windows, other |
| 5,760 | 2,643,274 | 0 | – |
| 5,840 | 981,136 | 3 | embedded |
| 16,000 | 781,225 | 5 | embedded |
| 4,096 | 775,473 | 5 | embedded |
| 1,024 | 758,230 | 4 | embedded |
| 2,800 | 677,211 | 1 | TP-Link router |

The good news is that the distribution of window size is heavily skewed toward well-known values, as seen in Table VIII. The most common window is unique to Linux variants, while the most ambiguous is split across 17 operating systems. Interestingly, window size 5760 in position #5, which we later discovered belongs to Ubuntu, is absent not just from ours, but also other fingerprinting databases (e.g., p0f, xprobe). We come back to these hosts later in the section and examine how Hershel classifies them. Ideally, unknown devices should be mapped to the same OS family (i.e., Linux in this case).

Another peculiar case is 168K hosts with zero window size, which in our database corresponds to a single device related to building automation. This particular stack forces the sender to finish the 3-packet handshake (SYN, SYN-ACK, ACK) and wait for the window to move before sending the first GET request. Immediately after the sender's ACK, the window expands to 12 288 B. Closed receiver windows can be an artifact of rate-limiting firewalls or site policies related to congestion control. One notable example is a popular host craigslist.com that prior to 2006 was completing all TCP handshakes with window size zero [25]. Other usage of this technique comes from network tarpits [2], which aim to slow down scanners by advertising small windows in SYN-ACKs. All of this suggests that the true window size may remain "hidden" from the fingerprinting tool for reasons unrelated to users, scrubbers, or TCP socket options.

The TTL values of received packets are plotted in Fig. 5(a), covering 251 unique points out of the 255 possible. A vast majority of the hosts are clustered on the values just before the initial TTL defaults 64, 128, and 255. Fig. 5(b) shows the distribution of reverse hop length for each host back to the scanner, calculated by subtracting the received TTL from the nearest well-known initial value. This distribution appears reasonable, with less than 0.4% of the mass below 10 or above 30 hops. This suggests the number of nonstandard initial TTLs (if any) is small. Table IX shows the distribution seen by Hershel and the corresponding number of signatures in $\mathcal{D}$.

TABLE IX
INITIAL TTL DISTRIBUTION (100% OF TOTAL)

| TTL | Hosts | Sigs | Group |
|---|---|---|---|
| 64 | 26,275,301 | 70 | linux, embedded, other |
| 128 | 7,129,667 | 17 | windows, embedded, other |
| 255 | 4,214,927 | 22 | linux, embedded, other |
| 32 | 190,697 | 7 | embedded |

TABLE X
BREAKDOWN OF 5.9 M HOSTS WITH RSTS

| Feature | Hosts | Sigs | Group |
|---|---|---|---|
| RA= 1, RN= 1 | 4,368,098 | 35 | embedded, other |
| RA= 0, RN= 1 | 1,167,761 | 11 | windows, embedded |
| RA= 0, RN= 0 | 367,915 | 10 | embedded |
| RA= 1, RN= 0 | 37,113 | 0 | – |

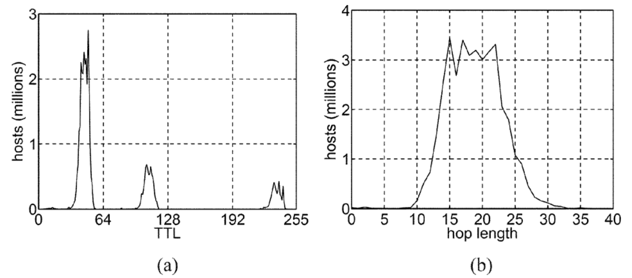


Fig. 5. Received TTL and reverse path length. (a) Received TTL. (b) Reverse distance.

A good number of hosts (69%) set the DF flag, indicating they intend to perform path-MTU discovery, which matches 45% of the signatures. Out of 37.8M responsive targets, 5.9M (16%) send at least one reset packet (in addition to the SYN-ACKs), which is consistent with 56 OSs. The reset window (RW) deviates from that in the SYN-ACK for 20.8% of the IPs and 8 fingerprints in $\mathcal{D}$.

Table X examines the interplay between RA and RN in reset packets. In the most common scenario, hosts indicate that the ACK sequence is valid and correctly acknowledge values one larger than transmitted by the scanner in the SYN packet (which encodes the destination IP). However, there are also 37K hosts (last row) with broken implementations that indicate a valid ACK, but set the field to zero. None of our signatures exhibits this behavior.

We have 21 unique combinations of options in $\mathcal{D}$. However, the dataset shows 264 different strings, with the top 10 provided in Table XI. Similar to Table VIII, a few popular cases account for the majority of IPs and Linux variants hold a clear lead, but now the most ambiguous combination splits across 41 embedded devices. While Akamai currently reports 137K servers [1], it seems reasonable that multiple NICs and IP aliasing can produce 339K samples in last row.

TABLE XI
TOP OPTIONS STRINGS (95% OF TOTAL)

| Options | Hosts | Sigs | Group |
|---|---|---|---|
| MSTNW | 13,156,171 | 8 | linux |
| MNWNNT | 6,214,837 | 18 | embedded, other |
| MNWNNTNNS | 5,579,866 | 12 | windows, other |
| M | 5,431,682 | 41 | embedded |
| MNW | 2,656,342 | 5 | linux, embedded, other |
| MNWST | 1,107,935 | 2 | windows, unix |
| MNWNNTSEE | 762,593 | 4 | other |
| MNNSWNNNT | 412,602 | 0 | – |
| MST | 370,699 | 1 | Windows Vista/7 |
| MNNSNW | 339,215 | 1 | Akamai linux |

TABLE XII
TOP MSS VALUES (93% OF TOTAL)

| MSS | Hosts | Sigs | Group |
|---|---|---|---|
| 1460 | 21,969,799 | 70 | all |
| 512 | 3,523,272 | 9 | embedded |
| 1452 | 3,512,626 | 2 | embedded |
| 1380 | 1,633,852 | 3 | windows, embedded |
| 1440 | 1,472,969 | 2 | linux, embedded |
| 1400 | 1,074,502 | 2 | embedded |
| 536 | 620,013 | 7 | embedded |
| 1448 | 562,961 | 0 | – |
| 1420 | 431,720 | 1 | Avocent KVM switch |
| 768 | 419,326 | 2 | embedded |

Practically every host (99.5%) supports the MSS option, with Table XII showing the top 10 cases out of the 1021 observed in the dataset. The most common MSS 1460 does not provide much information about the OS, but the other values appear useful at partitioning the dataset into small groups. On the downside, general-purpose OSs often set the MSS as a function of the underlying data-link layer (i.e., MSS = MTU − 40), which creates some interesting dilemmas. For example, MSS 1452 in third place can be classified as one of two embedded devices or as home computers with 1492-B MTUs commonly seen over PPP links such as DSL. This emphasizes importance of Hershel's probabilistic matching (15) and explains the significantly smaller number of unique MSS values in $\mathcal{D}$ (i.e., only 20).

## B. Classification Overview

We run Hershel on the scan dataset and obtain a nonzero classification probability for 37.4M devices. Before showing these results, we perform additional sanity checks by examining how often individual features of each IP matched those in the most-likely OS suggested by Hershel.

Starting with the first two columns of Table XIII, observe that window size is quite volatile, with 30% of the decisions going to signatures with a different window. This was expected given the numerous reasons to modify this field and the large amount of unique values seen earlier. Additionally, these 30% cover unknown devices whose RTOs and other features may match some OS in $\mathcal{D}$, but not the window size. Hershel remains robust in these cases and simply identifies the closest signature based on the available information. For example, 98.4% of Ubuntu cases

TABLE XIII
HERSHEL'S FEATURE MATCH RATE

| Feature | Fraction | RST possibilities | Fraction |
|---|---|---|---|
| Win | 70.3% | Neither has RST | 80.9% |
| TTL | 95.2% | Both have RST, match | 10.4% |
| DF | 96.2% | Missing RST | 4.2% |
| MSS | 70.6% | Both have RST, non-match | 3.5% |
| OPT | 99.4% | Bogus extra RST | 1.0% |

TABLE XIV
TOP INDIVIDUAL SIGNATURES (65% OF TOTAL)

| OS | Count |
|---|---|
| Linux 2.6 / 2.4 | 9,610,732 |
| VxWorks embedded systems | 4,179,583 |
| Windows Server 2003 SP1 SP2 | 2,316,590 |
| VxWorks 5.4.2 / Xerox embedded | 1,890,585 |
| Linux 2.6 / Debian / CentOS / SonicWall | 1,196,143 |
| Embedded Linux / Mikrotik routers | 1,190,102 |
| Windows Server 2008 SP1 SP2 R2 / Vista / 7 | 1,146,609 |
| TP-Link / Iball / Huawei home routers | 1,046,985 |
| Windows Server 2003 / 2000 / XP SP1 | 1,001,343 |
| Cisco / Scientific Atlanta cable modems | 827,285 |

with the unknown window 5760 are classified to Linux 2.4/2.6. These 2.6M hosts account for 25% of all window mismatch.

TTL and DF both exhibit match rates over 95%, while MSS comes in much lower at 71%. This is not surprising in light of its dependency on the MTU. The OPT string proves extremely reliable, where 77.4% of the cases match exactly and 22% are feasible subsets/supersets of the original. The five possible cases with RST packets are shown in the other two columns of Table XIII. Combining the first two rows, we can conclude that 91% of the hosts have a matching RST feature. The next row with missing RSTs allows us to ballpark network packet loss at $q_{\text{real}} = 4.2\%$, not too far from the model's 3.8%. The majority of nonmatching combinations (RA, RN, RW), responsible for 3.5% in the table, are caused by RW. Some of this behavior was also expected since tweaking of window size causes certain OSs to alter RW as well. Finally, we see 1% of the cases with extra RST packets, which we suspect are injected by firewalls, NAT boxes, and other devices as indication that they have expired the per-flow state.

## C. Results

Having verified the general soundness of Hershel's output, we show it in Table XIV. Linux attracts the most classification decisions, accounting for nearly a quarter of the Web servers. This signature is quite unique, which makes accidental lumping of unknown devices or misclassified hosts into this category highly unlikely. In second and fourth place is VxWorks, which is an embedded OS extensively used in routers, modems, cameras, and printers. Interestingly, Windows 2003 is third, well above Server 2008 in seventh position. More Linux, home routers/modems, and Server 2003/XP make up the remaining OSs.

Table XV groups fingerprints by type. Linux not just takes the first spot, but it dominates all other types of unix combined by a factor of 6. Embedded systems continue in second place, while windows is firmly in third. Interestingly, these results differ quite a bit from those in prior application of Snacktime

TABLE XV
COMMON FAMILIES OF OPERATING SYSTEMS

| Group | Count |
|---|---|
| Linux | 13,882,999 |
| Embedded | 13,590,803 |
| Windows | 7,561,839 |
| Other | 2,396,455 |

TABLE XVI
MANUAL VERIFICATION

| | IPs | Result | Count |
|---|---|---|---|
| Consensus | 429 | Both correct | 424 |
| | | Neither correct | 3 |
| | | Indeterminate | 2 |
| Disagreement | 571 | Hershel correct | 476 |
| | | Snacktime correct | 9 |
| | | Neither correct | 6 |
| | | Indeterminate | 80 |

to this dataset [24], with the most noticeable difference being 9M hosts moving from windows to embedded. This is not surprising as Snacktime's ability to overcome noise, packet loss, and feature corruption is quite weak. Furthermore, as shown above, Microsoft OSs often share the window size and TTL with embedded devices, making this distinction even more difficult for Snacktime.

To better understand the difference between these methods, we carry out comparison using manual analysis of 1000 random targets for which we had an HTTP header from a separate download process that grabbed the root page of each replying IP (this was done in real time during the 2010 scan). Table XVI shows the result. The first category in the table breaks down 429 hosts on which both methods produce the same exact OS. Out of these, 424 are correct matches, 3 incorrect, and 2 indeterminate. The last option occurs for devices inadequately represented in the database (i.e., no resemblance to any signature) or when multiple OSs appear to be probable (e.g., due to extensive packet loss or missing/ambiguous "Server:" field in the HTTP response header). Among the 571 disputed hosts, Hershel delivers 476 correct results and Snacktime 9.

We can make a decision for 918 cases, out of which Hershel's accuracy is 98% and Snacktime's is 47%. The 9 cases where Hershel is wrong, but Snacktime is right, are caused by bogus RSTs, which Snacktime ignores, but Hershel takes into account. Overall, we find that when the two methods disagree, Hershel is overwhelmingly more accurate.

### D. World View

Next, we use the MaxMind GeoIP database [13] to glean trends in OS usage around the globe. Table XVII shows the top countries in the measurement. The US leads the list, accounting for almost half of the discovered Web servers (i.e., 16M out of 37M) and exceeding China in second place by a factor of 8. The distribution of OS popularity is quite diverse, with only Italy and Brazil exhibiting similar vectors. Interestingly, Linux prevails over Windows in all countries except China; Spain stands out with 90% Linux, far more than any other locale in the list; and the US has the highest fraction of embedded devices among the entries in the table.

TABLE XVII
TOP COUNTRIES RUNNING WEB SERVERS (71% OF TOTAL)

| Country | Hosts | Windows | Linux | Embedded | Other |
|---|---|---|---|---|---|
| US | 16,187,542 | 16.1% | 30.4% | 47.2% | 5.2% |
| CN | 2,345,462 | 54.1% | 14.2% | 14.6% | 16.1% |
| ES | 1,620,920 | 4.1% | 89.2% | 5.5% | 0.8% |
| JP | 1,614,724 | 11.3% | 37.0% | 35.8% | 14.7% |
| DE | 1,043,699 | 19.9% | 57.7% | 15.7% | 5.7% |
| GB | 862,571 | 32.1% | 34.8% | 25.0% | 5.9% |
| CA | 849,285 | 25.9% | 45.3% | 12.8% | 14.7% |
| IT | 810,104 | 14.3% | 53.3% | 29.1% | 1.7% |
| BR | 685,597 | 14.5% | 52.8% | 25.2% | 5.3% |
| TW | 644,645 | 35.9% | 47.2% | 10.8% | 5.3% |

Table XVIII breaks down the data by AS, shedding additional light on the results. Home access providers in the US (i.e., Comcast, Time Warner, Cox) are full of embedded devices, likely consumer routers and modems. In combination, these 4.8M boxes represent 30% of the discovered servers in the US, which helps explain the high percentage of embedded stacks seen earlier. Similarly, Telefonica de Espana, a large telecommunications provider in Spain and South America, is responsible for 50% of Spanish Web servers in our dataset. This company is known for collaborations with RedHat and a cloud-computing emphasis [48]. Its 92% bias toward Linux is consistent with an earlier observation that Spain is dominated by this operating system. China's propensity toward Windows may stem from lax software piracy laws, with 67% of its devices coming from two ISPs in Table XVIII, each replete with Microsoft OSs.

### E. Scrubbers

While the Hershel's main purpose is large-scale measurement, where OS scrubbing is not likely to be prevalent, it still makes sense to examine its performance in such scenarios. Table XIX lists four obfuscators mentioned in existing literature and available for testing.

The first is Linux iptables, part of the packet-filtering framework called netfilter [29]. It is commonly used to inspect packets, modify routing tables, and configure the kernel firewall. It has extensions that "mangle" packets and change certain header fields. However, the only ones of interest to Hershel are TTL and MSS. OSfuscate [11] is a Windows scrubber that thwarts fingerprinting tools by changing the registry. It can modify Win, TTL, MSS, and certain options (i.e., drop SACK and timestamps). Along similar lines, TCP Optimizer [32] gives its users ability to change the same five registry values, in hopes of improving TCP transfer speed. Finally, IPPersonality [43], built on top of the netfilter framework, is the most sophisticated scrubber in the list. It can modify all Hershel features except RST and RTO.

To evaluate performance against scrubbers, we simulate the worst-case scenario—IPPersonality with an adversary who mimics the signature with the closest RTO vector from another OS family (i.e., windows, linux, embedded, other). Table XX shows the result using Pareto OWDs ($\mu = 0.5$ s) and the Zipf setup from Table VI. Snacktime stays in the single

TABLE XVIII
TOP ASS RUNNING WEB SERVERS (22% OF TOTAL)

| AS | Size | Owner | Hosts | Windows | Linux | Embedded | Other |
|---|---|---|---|---|---|---|---|
| 7922 | 71.0M | Comcast Cable | 3,444,634 | 3.3% | 6.2% | 89.8% | 0.3% |
| 4134 | 109.7M | Chinanet | 988,397 | 50.7% | 13.3% | 15.9% | 18.5% |
| 3352 | 10.9M | Telefonica de Espana | 861,222 | 2.3% | 92.0% | 4.8% | 0.5% |
| 4837 | 54.5M | CNC Group China | 595,931 | 53.2% | 9.4% | 15.0% | 21.7% |
| 20001 | 5.7M | Time Warner Cable | 485,766 | 2.4% | 1.5% | 95.4% | 0.3% |
| 11351 | 4.9M | Time Warner Cable | 436,329 | 2.0% | 1.1% | 96.3% | 0.2% |
| 2914 | 7.7M | NTT America | 429,648 | 25.6% | 20.6% | 20.3% | 33.1% |
| 22773 | 11.9M | Cox Communications Inc. | 426,807 | 4.8% | 2.8% | 90.8% | 0.6% |
| 7018 | 75.2M | AT&T Services | 373,068 | 31.0% | 36.2% | 18.9% | 11.0% |
| 7155 | 988K | Viasat Communications Inc. | 370,821 | 39.0% | 0.0% | 60.9% | 0.0% |

TABLE XIX
CAPABILITY OF OS OBFUSCATION TOOLS

| Tool | Win | TTL | DF | TCP Options | MSS |
|---|---|---|---|---|---|
| Netfilter iptables | – | X | – | – | X |
| OSfuscate | X | X | – | Drop S and T | X |
| SG TCP Optimizer | X | X | – | Drop S and T | X |
| IPPersonality | X | X | X | Replace or reorder | X |

TABLE XX
SCRUBBED ACCURACY (PERCENT) AMONG ALL OSS

| Loss | Snacktime | | Hershel | | | | |
|---|---|---|---|---|---|---|---|
| (%) | All | RTO | -RST-RTO | -RST | RTO | All | RTO+RST |
| 0.0 | 9.9 | 12.6 | 0.0 | 11.8 | 22.1 | 36.6 | 47.6 |
| 3.8 | 7.8 | 10.0 | 0.0 | 11.4 | 21.4 | 34.8 | 45.3 |
| 10 | 5.2 | 6.9 | 0.0 | 10.9 | 20.1 | 31.9 | 41.8 |
| 50 | 0.6 | 0.8 | 0.0 | 6.0 | 10.4 | 15.6 | 20.5 |

TABLE XXI
SCRUBBED ACCURACY (PERCENT) AMONG WINDOWS/LINUX

| Loss | Snacktime | | Hershel | | | | |
|---|---|---|---|---|---|---|---|
| (%) | All | RTO | -RST-RTO | -RST | RTO | All | RST+RTO |
| 0.0 | 2.5 | 5.9 | 0.0 | 5.0 | 14.6 | 31.8 | 41.7 |
| 3.8 | 2.0 | 5.0 | 0.0 | 4.7 | 14.1 | 29.8 | 39.4 |
| 10 | 1.4 | 3.8 | 0.0 | 4.5 | 12.9 | 26.9 | 35.7 |
| 50 | 0.1 | 0.5 | 0.0 | 2.2 | 6.2 | 12.0 | 16.3 |

digits, showing performance slightly below that of using just the RTOs. Hershel with only the fixed features from previous literature (i.e., all except RTO and RST) produces the expected 0% match rate. Adding the RTO pushes accuracy to 6%–12%, but this far from impressive—the RTO alone works better, achieving 10%–22%. Employing all Hershel features almost doubles the result. However, the real winner in this comparison is the RST + RTO combination, which reaches as high as 47%.

Limiting the simulation to 26 Windows/Linux signatures that the scrubber modifies using the same rules produces a more challenging case outlined in Table XXI. There is an accuracy reduction in all categories, but the scrubber-resilient version of Hershel still manages to correctly pinpoint over 41% of the samples that sustain no loss.

## VII. CONCLUSION

We modeled the problem of single-packet OS fingerprinting and developed novel approaches for tackling delay jitter, packet loss, and user modification to SYN-ACK features. Based on this theory, we created a classification method that significantly increased the accuracy of existing techniques, both in simulation and the real Internet.

Future work involves multipass extraction of jitter, packet-loss, and OS-popularity models from the observed samples, which should improve estimation accuracy even further.

## REFERENCES


[1] Akamai, "Akamai," [Online]. Available: http://www.akamai.com/html/about/facts_figures.html

[2] L. Alt, R. Beverly, and A. Dainotti, "Uncovering network tarpits with degreaser," in *Proc. ACM ACSAC*, Dec. 2014, pp. 156–165.

[3] P. Auffret, "SinFP, unification of active and passive operating system fingerprinting," *J. Comput. Virology*, vol. 6, no. 3, pp. 197–205, Nov. 2010.

[4] T. Beardsley, "Snacktime: A Perl solution for remote OS fingerprinting," Jun. 2003 [Online]. Available: http://www.planb-security.net/wp/snacktime.html

[5] R. Beck, "Passive-aggressive resistance: OS fingerprint evasion," *Linux J.*, vol. 2001, no. 89, Aug. 2001.

[6] D. B. Berrueta, "A practical approach for defeating Nmap OS-fingerprinting," 2003 [Online]. Available: http://nmap.org/misc/defeat-nmap-osdetect.html

[7] R. Beverly, "A robust classifier for passive TCP/IP fingerprinting," in *Proc. PAM*, Apr. 2004, p. 158.

[8] R. Braden, "Requirements for Internet hosts—Communication layers," IETF RFC 1122, Oct. 1989.

[9] J. Caballero *et al.*, "FiG: automatic fingerprint generation," in *Proc. NDSS*, Feb. 2007, pp. 27–42.

[10] H. K. J. Chu, "Tuning TCP parameters for the 21st century," Jul. 2009 [Online]. Available: http://www.ietf.org/proceedings/75/slides/tcpm-1.pdf

[11] A. Crenshaw, "OSfuscate," 2008 [Online]. Available: http://www.irongeek.com/i.php?page=security/code

[12] D. Dagon, N. Provos, C. P. Lee, and W. Lee, "Corrupted DNS resolution paths: the rise of a malicious resolution authority," in *Proc. NDSS*, Feb. 2008.

[13] "Maxmind GeoIP databases," [Online]. Available: http://dev.maxmind.com/geoip/

[14] T. Dunigan, M. Mathis, and B. Tierney, "A TCP tuning daemon," in *Proc. ACM/IEEE Supercomput.*, Nov. 2002, pp. 1–16.

[15] Z. Durumeric, E. Wustrow, and J. Halderman, "ZMap: Fast Internet-wide scanning and its security applications," in *Proc. USENIX Security*, Aug. 2013, pp. 605–620.

[16] L. G. Greenwald and T. J. Thomas, "Toward undetected operating system fingerprinting," in *Proc. USENIX WOOT*, Aug. 2007, pp. 1–10.

[17] S. Guoqiang and D. Lee, "Network protocol system fingerprinting: a formal approach," in *Proc. IEEE INFOCOM*, Apr. 2006, pp. 1–12.

[18] J. Heidemann *et al.*, "Census and survey of the visible Internet," in *Proc. ACM IMC*, Oct. 2008, pp. 169–182.

[19] M. Honda *et al.*, "Is it still possible to extend TCP?," in *Proc. ACM IMC*, Nov. 2011, pp. 181–194.

[20] "IRL fingerprinting dataset," [Online]. Available: http://irl.cs.tamu.edu/projects/sampling/

[21] V. Jacobson, R. Braden, and D. Borman, "TCP extensions for high performance," IETF RFC 1323, May 1992.

[22] T. Kohno, A. Broido, and K. C. Claffy, "Remote physical device fingerprinting," *IEEE Trans. Depend. Secure Comput.*, vol. 2, no. 2, pp. 93–108, May 2005.
[23] E. Kollmann, "Chatter on the wire: A look at DHCP traffic," 2007 [Online]. Available: http://chatteronthewire.org/download/chatter-dhcp.pdf
[24] D. Leonard and D. Loguinov, "Demystifying service discovery: implementing an Internet-wide scanner," in *Proc. ACM IMC*, Nov. 2010, pp. 109–122.
[25] J. Lippard, "Craigslist no longer uses TCP window size of 0," 2006 [Online]. Available: http://lippard.blogspot.com/2006/07/craigslist-no-longer-uses-tcp-window.html
[26] J. Medeiros, A. Brito, and P. Pires, "A data mining based analysis of nmap operating system fingerprint database," in *Proc. IEEE CISIS*, Sep. 2009, pp. 1–8.
[27] J. Medeiros, A. Brito, and P. Pires, "An effective TCP/IP fingerprinting technique based on strange attractors classification," in *Proc. DPM/SETOP*, Sep. 2009, pp. 208–221.
[28] "Microsoft support," [Online]. Available: http://support.microsoft.com/kb/2525390
[29] D. Napier, "IPTables/NetFilter—Linux's next generation stateful packet filter," *SysAdmin Mag.*, vol. 10, pp. 8–16, Nov. 2001.
[30] "Netcraft Web server survey," [Online]. Available: http://news.netcraft.com/
[31] "Nmap," [Online]. Available: http://nmap.org/
[32] "SpeedGuide TCP optimizer," [Online]. Available: http://www.speedguide.net/downloads.php
[33] Oracle, "Operating system tuning," [Online]. Available: http://docs.oracle.com/cd/E12839_01/web.1111/e13814/os_tuning.htm
[34] J. Padhye and S. Floyd, "On inferring TCP behavior," in *Proc. ACM SIGCOMM*, Aug. 2001, pp. 287–298.
[35] V. Paxson, M. Allman, J. Chu, and M. Sargent, "Computing TCP's retransmission timer," IETF RFC 6298, Jun. 2011.
[36] J. Postel, "Transmission control protocol," IETF RFC 793, Sep. 1981.
[37] G. Prigent, F. Vichot, and F. Harrouet, "IpMorph: fingerprinting spoofing unification," *J. Comput. Virology*, vol. 6, no. 4, pp. 329–342, Nov. 2010.
[38] N. Provos, "A virtual honeypot framework," in *Proc. USENIX Security*, Aug. 2004, pp. 1–14.
[39] N. Provos and P. Honeyman, "ScanSSH—Scanning the Internet for SSH servers," in *Proc. USENIX LISA*, Dec. 2001, pp. 25–30.
[40] Y. Pryadkin, R. Lindell, J. Bannister, and R. Govindan, "An empirical evaluation of IP address space occupancy," USC/ISI, Tech. Rep. ISI-TR-2004-598, 2004.
[41] D. Richardson, S. Gribble, and T. Kohno, "The limits of automatic OS fingerprint generation," in *Proc. ACM AISec*, Oct. 2010, pp. 24–34.
[42] M. Roesch, "Snort—Lightweight intrusion detection for networks," in *Proc. USENIX LISA*, Nov. 1999, pp. 229–238.
[43] G. Roualland and J.-M. Saffroy, "IP Personality," [Online]. Available: http://ippersonality.sourceforge.net/
[44] C. Sarraute and J. Burroni, "Using neural networks to improve classical operating system fingerprinting techniques," *Electron. J. SADIO*, vol. 8, no. 1, Mar. 2008.
[45] S. Shah, "An introduction to HTTP fingerprinting," May 2004 [Online]. Available: http://net-square.com/httprint_paper.html
[46] M. Smart, G. R. Malan, and F. Jahanian, "Defeating TCP/IP stack fingerprinting," in *Proc. USENIX Security*, Jun. 2000, pp. 229–240.
[47] "Snort IDS," [Online]. Available: http://www.snort.org
[48] "RedHat customer case study," 2014 [Online]. Available: https://www.redhat.com/en/files/resources/en-rh-telefonica-global-solutions-chooses-red-hat-plan-cloud-future-11772377.pdf
[49] G. Taleck, "Ambiguity resolution via passive OS fingerprinting," in *Proc. RAID*, Sep. 2003, pp. 192–206.
[50] G. Taleck, "SYNSCAN: Towards complete TCP/IP fingerprinting," presented at the CanSecWest, Apr. 2004.
[51] B. Tierney, "TCP tuning guide for distributed applications on wide area networks," *USENIX & SAGE Login*, vol. 26, no. 1, pp. 33–39, Feb. 2001.
[52] C. Valli, "Honeyd—A OS fingerprinting artifice," in *Proc. Australian Comput., Netw. Inf. Forensics Conf.*, Nov. 2003.
[53] F. Veysset, O. Courtay, O. Heen, and I. R. Team, "New tool and technique for remote operating system fingerprinting," Apr. 2002 [Online]. Available: http://www.ouah.org/ring-full-paper.pdf
[54] K. Wang, "Frustrating OS fingerprinting with morph," 2004 [Online]. Available: https://defcon.org/images/defcon-12/dc-12-presentations/Wang/dc-12-wang.pdf
[55] F. V. Yarochkin *et al.*, "Xprobe2++: Low volume remote network information gathering tool," in *Proc. IEEE/IFIP DSN*, Jun. 2009, pp. 205–210.
[56] M. Zalewski, "Strange attractors and TCP/IP sequence number analysis," Apr. 2001 [Online]. Available: http://lcamtuf.coredump.cx/newtcp/
[57] M. Zalewski, "P0f v3: Passive fingerprinter," 2012 [Online]. Available: http://lcamtuf.coredump.cx/p0f3
[58] Y. G. Zeng, D. Coffey, and J. Viega, "How vulnerable are unprotected machines on the Internet?," in *Proc. PAM*, Mar. 2014, pp. 224–234.

**Zain Shamsi** received the B.S. degree (with honors) in computer science from Trinity University, San Antonio, TX, USA, in 2008, and is currently pursuing the Ph.D. degree in computer science from Texas A&M University, College Station, TX, USA.

His research interests include Internet measurement, network protocols, and security.

**Ankur Nandwani** received the B.S. degree in computer engineering from the National Institute of Technology, Surat, India, in 2008, and the M.S. degree in computer science from Texas A&M University, College Station, TX, USA, in 2010.

He currently leads the Mobile Development Team with Coinbase, Inc., San Francisco, CA, USA. His research interests include mobile security, machine learning, and crypto-currencies.

**Derek Leonard** (S'05) received the B.A. degree (with distinction) in computer science and mathematics from Hendrix College, Conway, AR, USA, in 2002, and the Ph.D. degree in computer science from Texas A&M University, College Station, TX, USA, in 2010.

He is currently a Data Scientist with Acxiom Corporation, Little Rock, AR, USA. His research interests include large-scale measurement of the Internet and peer-to-peer networks.

**Dmitri Loguinov** (S'99–M'03–SM'08) received the B.S. degree (with honors) from Moscow State University, Moscow, Russia, in 1995, and the Ph.D. degree from the City University of New York, New York, NY, USA, in 2002, both in computer science.

He is currently a Professor with the Department of Computer Science and Engineering, Texas A&M University, College Station, TX, USA. His research interests include P2P networks, information retrieval, congestion control, Internet measurement, and modeling.