

Faults: A Non-Parametric Iterative Classifier for Internet-Wide OS Fingerprinting

Zain Shamsi¹, Daren B. H. Cline², and Dmitri Loguinov³, *Senior Member, IEEE*

Abstract—Recent work in OS fingerprinting has focused on overcoming random distortion in network and user features during Internet-scale SYN scans. These classification techniques work under an assumption that all parameters of the profiled network are known a-priori – the likelihood of packet loss, the popularity of each OS, the distribution of network delay, and the probability of user modification to each default TCP/IP header value. However, it is currently unclear how to obtain realistic versions of these parameters for the public Internet and/or customize them to a particular network being analyzed. To address this issue, we derive a non-parametric Expectation-Maximization (EM) estimator, which we call *Faults*, for the unknown distributions involved in single-probe OS fingerprinting and demonstrate its significantly higher robustness to noise compared to methods in prior work. We apply *Faults* to a new scan of 67M web servers and discuss its findings.

Index Terms—Stack fingerprinting, network security.

I. INTRODUCTION

THE Internet is a fascinating conglomerate of highly heterogeneous devices, which differ in hardware capability, security awareness, software features, and daily usage. Measuring the amount, type, and behavior of these devices, as well the networks they connect to, has become an important topic [14], [15], [17], [20], [27], [30], [36], [45], [46]. To categorize the makeup of today’s networks, research in *active OS fingerprinting*, which is our topic in this paper, aims to determine the stack of remote hosts using their responses to external stimuli (i.e., TCP/IP probes) [4], [5], [7], [10], [18], [25], [26], [31], [33], [39], [44], [49], [53], [54], [58], [59], [60]. In addition to uncovering the operating system of computers, fingerprinting can expose household items (e.g., printers, cameras, TVs) and various cyber-physical systems (e.g., temperature monitors, lighting controllers), which are classes of devices that have enjoyed increased exploitation in recent years [2], [9].

There are many uses for remote stack fingerprinting. First, it helps hackers in identification of vulnerable hosts and general network reconnaissance [52], especially during

cyber-attacks that target only a specific OS implementation [22]. Second, OS fingerprinting is routinely deployed in security, e.g., by administrators of large networks seeking to find unpatched hosts and rogue entities [1], [32], [50]. Third, perimeter-defense systems (e.g., IDS, firewalls) may require the OS of the target host in order to detect certain types of exploits (e.g., those involving reassembly of IP fragments). In such cases, autonomous fingerprinting of the protected network allows these installations to function at maximum effectiveness [49], [53]. Finally, researchers/organizations use these techniques to understand usage trends [37], [38], discover the spread of new technologies [8], [16], [29], [41], and expose botnets [28].

Active stack fingerprinting can be partitioned into three categories – *banner-grabbing* via plain-text protocols (e.g., telnet, HTTP, FTP) [43], [57], *multi-probe* tools that elicit OS-specific responses from various non-standardized combinations of flags and/or unexpected usage of protocol fields (e.g., nmap [39], xprobe [58], p0f [60]), and *single-probe* methods that send a regular SYN to each host (e.g., Snacktime [6], RING [55], Hershel [46], Hershel+ [45]).

At large scale, banner-grabbing has several impediments – frequent removal of OS-identifying strings by administrators (e.g., for security purposes), high bandwidth overhead, and common interaction with non-platform-specific software (e.g., apache, nginx). Multi-probe tools have their own challenges – heavy load on the target, massive complaints about intrusive activity, and noisy results when the destination IP is load-balanced across a server farm (i.e., each packet hits a different machine). More importantly, the accuracy of multi-packet tools suffers a significant degradation when firewalls block auxiliary probes (e.g., a UDP to a closed port, rainbow flags in TCP headers, ICMP port unreachable) and the underlying classifier is not robust against unexpected feature removal/modification. As shown in [45], OS classification with nmap over the public Internet fails in almost 30% of the cases. Furthermore, nmap sometimes produces nonsensical results and worse accuracy than the alternatives utilizing a single probe.

Before modeling and improving multi-packet classifiers, which are still poorly understood, it is important to ask whether there exists a set of algorithms for maximizing performance of single-packet tools in Internet-wide scans. Such techniques provide a maximally stealthy option and may be able to bypass firewalls/IDS when packets loaded with “tricks” cannot. As it turns out, even the most advanced model in single-probe literature, i.e., Hershel+ [45], leaves room for improvement. It has many built-in assumptions that may be

Manuscript received October 29, 2019; revised July 10, 2020 and January 14, 2021; accepted June 4, 2021; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor E. Kalyvianaki. Date of publication June 18, 2021; date of current version October 15, 2021. (*Corresponding author: Dmitri Loguinov.*)

Zain Shamsi is with Applied Research Laboratories (ARL), Austin, TX, USA (e-mail: zain.shamsi@utexas.edu).

Daren B. H. Cline is with the Department of Statistics, Texas A&M University, College Station, TX 77843 USA (e-mail: dcline@stat.tamu.edu).

Dmitri Loguinov is with the Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77843 USA (e-mail: dmitri@cse.tamu.edu).

Digital Object Identifier 10.1109/TNET.2021.3088333

violated in practice, which in turn may affect its classification accuracy and overall performance on such basic metrics as the fraction of the Internet running a particular stack. Our motivation for this paper is to understand the limitations of existing single-probe techniques and offer novel avenues for increasing both the classification accuracy and amount of information recovered from responses to a SYN packet.

Compared to the conference version [47], this paper includes a revised explanation of the material, more challenging distortion scenarios (i.e., S_{14} in Tables II-IV and S_{24} in Tables V-VI), Figs. 4(c)-(d) with additional accuracy verification of the proposed estimator, and Section VI.D that focuses on classification of unknown signatures.

II. BACKGROUND

A. Nmap

Perhaps the most popular and exhaustive tool for OS fingerprinting is nmap [39]. To understand its infeasibility for wide-area usage, we briefly review its outgoing traffic and response requirements, as well as the matching algorithm. By default, nmap starts with a vertical scan of the target using 1,000 well-known ports in an attempt to find two TCP ports, one of which is open and the other is closed, as well as a closed UDP port. It then sends 16 uniquely crafted probes – six regular TCP packets to an open port, one valid and one invalid ICMP ping, one UDP packet to a closed port, four malformed packets to an open TCP port, and three TCP packets to a closed port. It retransmits all probes multiple times to neutralize the impact of packet loss, which results in over 100 packets per host in addition to the initial port scan.

Besides overhead, running nmap against the entire Internet poses a number of additional challenges. First, there is a low likelihood that a port scan, combined with probes to closed ports, gets unnoticed by the IDS. Many software packages (e.g., snort) contain explicit rules to detect and block the rather esoteric nmap traffic. Certain networks take offense at being nmapped, which results in swift action to block the entire subnet/AS of the scanner and complaints about abusive behavior. Second, the firewall may allow select ports to reach the target host (e.g., port 80 to a webserver); however, there is little incentive to pass UDP or TCP packets to other ports that do not offer any services. Third, in similar fashion, the OS firewall can be configured (e.g., using domain group policy) to silently drop incoming packets to closed ports. In fact, Windows and Mac OS X suppress outgoing ICMP port unreachables *even when an explicit rule is created to allow such packets through the firewall* [3], [34].

Nmap expects responses to not deviate from those specified in the database (e.g., a RST to a TCP rainbow packet, ICMP port unreachable from a closed UDP port, ICMP echo reply to a ping). Because it considers absence of a response to be a feature, it can be misled into assigning large positive weights to firewall actions, which skews the result towards network stacks that inherently respond with fewer signals. This may occur despite a complete non-match in other features, meaning that the target may share nothing in the packet header (e.g., TCP window size, TTL, options, MSS) with the signature it is matched to [45]. Other issues include the database itself, which

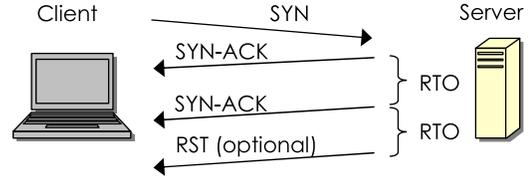


Fig. 1. Half-open connections in TCP.

contains signatures that are subsets of others from completely unrelated stacks and allows special *null* header fields that can match any value in the observation. Unless the target responds to all 16 probes exactly as expected, an obscure device with the most null fields can trump the other alternatives, including the correct signature.

Additionally, certain TCP fields are quite *volatile*, i.e., change from user tweaks, underlying network MTU, and software setsockopt function calls. This does not inherently change the operating system, but creates an illusion of a different stack. For example, Server 2008 R2 accepts incoming connections with a kernel buffer (i.e., TCP window size) of 8,192 bytes; however, an apache webserver can reconfigure this field to an arbitrary value before listening on the socket. Furthermore, this can be done on a per-socket basis and may vary over time depending on memory usage or other considerations. When faced with this type of uncertainty, nmap uses heuristic weights and thresholds that do not have rigorous theory/verification behind them. As a result, it exhibits highly unreliable identification in certain scenarios [45].

B. Single-Packet Tools

For accurate OS fingerprinting at Internet scale, low-overhead methods resilient to volatility are preferred. Our focus in this paper is on single-probe techniques, which generally work by sending a TCP SYN to the target host and inducing a stream of SYN-ACK responses, possibly with a RST at the end. Since the connection is kept in the half-open state, the server continues retransmitting SYN-ACKs until its internal maximum-retry threshold is exceeded. Delays between the SYN-ACKs, known as *retransmission timeouts* (RTOs), as well as their count and presence of the last RST, reveal valuable information about the OS of the responding host. This is illustrated in Fig. 1. Coupling the RTOs with default TCP/IP header values makes stack classification possible.

The main difference between prior work [6], [27], [45], [46], [55] lies in the features they extract from TCP/IP headers and the assumed distortion model. As of this writing, Hershel+ [45] is both the most recent effort in this direction and most robust to observation noise. We review its operation and formulas later in the paper.

III. LEARNING FROM OBSERVATION

A. General Problem

Suppose the OS database consists of $n \geq 1$ known stacks $(\omega_1, \dots, \omega_n)$, each with some vector-valued *fingerprint* $\mathbf{x}_i = (x_{i1}, x_{i2}, \dots)$. Table I shows three examples from the Plata database [45], where each row represents an OS and its fingerprint. Vector \mathbf{x}_i contains a combination of

TABLE I
FEATURE VECTORS \mathbf{x}_i (TCP OPTIONS: M = MSS, N = NOP, S = SACK, T = TIMESTAMP, AND W = WINDOW SCALE)

OS name	Win	TTL	DF	OPT	MSS	RST	RTOs (sec)
Linux 3.2	5,792	64	1	MSTNW	1,460	0, 0, 0, 0	3, 6, 12, 24.2, 48.2
Windows 2003	16,384	128	0	MNWNNTNNS	1,380	0, 0, 0, 0	3, 6.5
Novell	6,144	128	1	MNWSNN	1,460	1, 1, 0, 1	1.4, 3.0

fixed features that are sent by the OS for new connections (i.e., TCP receiver window size, TTL, the do-not-fragment flag, supported TCP options, MSS, RST peculiarities) and SYN-ACK retransmission timeouts (RTOs). Further assume a set of observations $\mathbf{x}' = (\mathbf{x}'_1, \dots, \mathbf{x}'_m)$ obtained by scanning the Internet and eliciting responses from m live servers, where $\mathbf{x}'_j = (x'_{j1}, x'_{j2}, \dots)$ is a vector of sampled features from host j . For the type of OS fingerprinting considered here, i.e., single-probe, this is done by dispatching a SYN to every IP address in BGP and collecting SYN-ACKs/RSTs from the responding servers, as previously shown in Fig. 1.

The goal of the classifier is then to determine for each \mathbf{x}'_j the most-likely fingerprint in the database. This task is complicated by the presence of distortion (also called *volatility* [45]) θ that randomly modifies the original features of the system before the observer gets them. This may involve a change in the temporal relationship between the packets (e.g., queuing delays), removal of some features (e.g., loss of RST packets), and rewriting of TCP headers in an effort to optimize or obscure the end-system.

Define $\alpha_i = p(\omega_i)$ to be the unknown fraction of hosts in \mathbf{x}' with OS i and let $\alpha = (\alpha_1, \dots, \alpha_n)$ be the corresponding vector. Now suppose $p(\mathbf{y}|\omega_i, \theta)$ is the probability that the fingerprint of signature i has been changed into \mathbf{y} under θ . Similarly, assume that $p(\omega_i|\mathbf{y}, \theta, \alpha)$ is the probability that an observed vector \mathbf{y} was produced by a host running OS i , conditioned on distortion model θ and popularity α . Then, application of Bayes' rule shows that the classifier must determine for each j the one database entry ω_i with the largest

$$p(\omega_i|\mathbf{x}'_j, \theta, \alpha) = \frac{\alpha_i p(\mathbf{x}'_j|\omega_i, \theta)}{p(\mathbf{x}'_j|\theta, \alpha)}, \quad (1)$$

where, for any vector of features \mathbf{y} , the denominator is

$$p(\mathbf{y}|\theta, \alpha) = \sum_{i=1}^n \alpha_i p(\mathbf{y}|\omega_i, \theta). \quad (2)$$

Analysis of (1) in existing work [45], [46] assumes that α is uniform (i.e., $\alpha_i = 1/n$) and θ is fixed by oracle input. Therefore, both α_i and denominator $p(\mathbf{x}'_j|\theta, \alpha)$ are independent of i and can be removed from the optimization, leaving only $p(\mathbf{x}'_j|\omega_i, \theta)$. In contrast, our goal here is to estimate both α and θ dynamically as the classifier is running, which should both increase its accuracy and yield interesting Internet-characterization results as byproduct of classification. Before reaching this objective, a gradual build-up of formalization is needed. This section deals with estimating α , the next one covers network distortion, and the one after that focuses on modification to fixed header features.

B. Fingerprint Popularity

Observation vector \mathbf{x}' gives rise to a number of equations in the form of (2), where the left side contains the empirical

(known) probability of observing each unique vector $\mathbf{y} \in \mathbf{x}'$ and the right side is a model that embeds the unknown parameters. Extraction of α and θ from such systems of equations commonly involves the Expectation-Maximization (EM) method, which produces a solution using fixed-point iteration [13], [19]. At every step t , it maximizes the expected log-likelihood function conditioned on the parameters obtained during the previous iteration $t - 1$. As long as the number of equations exceeds the number of unknown parameters, EM works well for many problems in practice.

For now, we treat $p(\mathbf{x}'_j|\omega_i, \theta)$ as a black-box classifier (e.g., Snacktime, Hershel, Hershel+), which does not attempt to estimate θ , and focus on determining α . This is the simplest (and only) case where (2) forms a linear system of equations, i.e., $p(\mathbf{y}_k|\theta, \alpha) = \sum_{i=1}^n \alpha_i c_{ik}(\theta)$, where all $c_{ik}(\theta)$ are constants. Throughout the paper, superscripts applied to parameters refer to the iteration number during which they are estimated, e.g., α_i^t approximates α_i during step t . Now notice that a sensible estimate of popularity for OS i is the average probability with which observations map to this fingerprint, conditioned on the previous estimate of popularity, i.e.,

$$\alpha_i^{t+1} = \frac{1}{m} \sum_{j=1}^m p(\omega_i|\mathbf{x}'_j, \theta, \alpha^t). \quad (3)$$

While the next result is fairly straightforward, its derivation methodology is needed for later parts of the paper.

Theorem 1: For a classifier with fixed θ , (3) represents the EM algorithm for recovering the popularity vector α .

Note that this is markedly different from deciding popularity using the fraction of classification decisions that go to each OS, which is known as *hard EM* and commonly used in clustering algorithms such as k -means [23]. In fact, all previous fingerprinting tools [6], [7], [39], [45], [46], [58], [60] can be viewed as performing one iteration of hard EM, i.e., outputting the fraction of classifications that belong to each OS ω_i as an estimate of its popularity α_i .

C. Discussion

We now address the question of whether (3) is sufficient for achieving good classification on its own and how much of the accuracy depends on knowing the exact distortion model θ . If the majority of the benefit is already obtained from recovering α , the extra computational cost and modeling effort involved in estimating θ may be unnecessary. For discussion purposes, we use a set of toy databases that allow simple demonstration of the intended effects. We apply a variety of distortions θ (from mild to radical) to demonstrate that a) prior methods can produce arbitrarily inaccurate estimation; b) both network and user distortion must be taken into account; and c) the new approach can recover the true underlying parameters of the network/user despite the heavy noise added

TABLE II
NETWORK DISTORTION IN SCENARIO S_1

Case	Forward latency (sec)		One-way delay (sec)		Loss
	Distribution	Mean	Distribution	Mean	
S_{11}	Erlang(2)	0.5	Exp	0.5	3.8%
S_{12}	Pareto	0.5	Pareto	0.5	50%
S_{13}	Reverse-exp	1.5	Erlang(2)	0.5	10%
S_{14}	Pareto	0.1	Uniform	0.1	50%

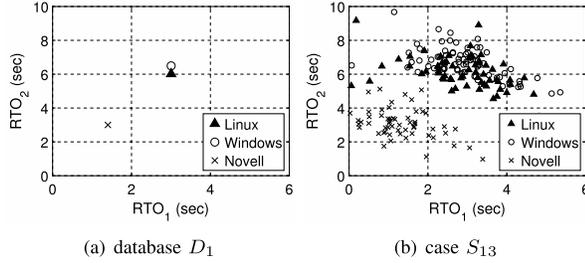


Fig. 2. Database and distorted observations.

to the observations. Note that the same observations arise with larger databases, but finding the corresponding scenarios may be more time-consuming.

Simulations below apply a forward latency to the SYN packet, pass each SYN-ACK through a FIFO queue, which adds random one-way delays along the return path, and drop packets using an iid (independent and identically distributed) loss model with some fixed probability. This is similar to the context in which prior methods [45], [46] have been tested. For the scenario we call S_1 , there are four different cases for the distribution of forward/reverse delays and packet-loss probability. These are shown in Table II and discussed next.

The first row matches exactly the assumed parameters θ in Hershel+ [45]. The second row uses Pareto delays with mean 500 ms and 50% loss, emulating highly noisy network conditions. The next row uses a shifted *reverse-exponential* forward latency with CDF $e^{-\lambda(2-x)}$, defined for $-\infty < x \leq 2$, which tests contrary-to-intuition examples where larger delays are more likely than smaller. We employ $\lambda = 2$ and truncate this distribution at zero, obtaining the average forward SYN delay of 1.5 sec. The last case in the table examines smaller average delays than the assumed model θ in Hershel+, but couples it with substantial loss.

Our first database D_1 contains truncated signatures of Linux 3.2 (ω_1), Windows Server 2003 (ω_2), and Novell Netware (ω_3) from Table I. We retain the first two retransmission timeouts (RTOs), remove all fixed header features, and obtain the fingerprints in Fig. 2(a). Note that these Linux and Windows signatures are pretty close to each other, albeit not identical; however, they are quite different from Novell. Distortion from case S_{13} applied to this database is illustrated in subfigure b), where we show the first 200 samples and remove observations with lost packets.

Define ρ^t to be the fraction of correct classifications for a given method during iteration t , where $t = \infty$ represents the convergence point of the underlying estimator (usually 20–40 iterations). If the method does not perform iteration, only ρ^1 is meaningful. We consider three techniques – Hershel+, hard EM with multiple iterations, and soft EM in (3), all

TABLE III
CLASSIFICATION RESULTS IN D_1

Case	α	Hershel+		Hard EM		EM in (3)	
		ρ^1	α^1	ρ^∞	α^∞	ρ^∞	α^∞
S_{11}	0.90	0.67	0.59	0.95	0.95	0.95	0.89
	0.05		0.35		0.00		0.06
	0.05		0.06		0.05		0.05
S_{12}	0.05	0.48	0.45	0.06	0.98	0.89	0.11
	0.90		0.41		0.00		0.82
	0.05		0.12		0.02		0.07
S_{13}	0.90	0.45	0.37	0.09	0.01	0.10	0.11
	0.05		0.51		0.88		0.79
	0.05		0.12		0.11		0.10
S_{14}	0.3	0.60	0.65	0.33	0.97	0.34	0.81
	0.6		0.23		0.00		0.13
	0.1		0.12		0.03		0.05

using the same function $p(\mathbf{x}'_j|\omega_i, \theta)$ and starting from uniform popularity $\alpha_i^0 = 1/n$.

Results of this process with $m = 2^{18}$ observations are shown in Table III. In the first row, Hershel+ performs quite well, achieving $\rho^1 = 67\%$. Since Novell Netware is an easy-to-separate signature from the other two, Hershel+ recovers α_3 pretty accurately; however, it is utterly confused about the frequency of the other two stacks. Applying hard EM increases accuracy, but full reconstruction of α still proves difficult. Application of (3) solves this issue.

Swapping (α_1, α_2) , the second simulation in Table III shows that Hershel+ is essentially guessing between Linux and Windows, while hard EM is misled into divergence, where it drops accuracy from 48% to 6%. While (3) is immune to divergence in this case, its estimate of α suffers from non-negligible errors. The next two cases in the table are even more difficult. They show that EM can be driven into inferior states when the assumed θ greatly deviates from that of the underlying network. In fact, application of (3) not only fails to obtain vectors α that resemble the true distribution, but also harms performance of the system, i.e., $\rho^\infty \ll \rho^1$.

IV. NETWORK FEATURES

A. Distortion Model

Our goal in this section is to estimate unknown distortion parameters θ inside $p(\mathbf{x}'_j|\omega_i, \theta)$. Let features $\mathbf{x}_i = (\mathbf{d}_i, \mathbf{u}_i)$ consist of network components (i.e., delays \mathbf{d}_i) and user-modified header values (i.e., \mathbf{u}_i). Since classification [45], [46] usually assumes that distortion is applied to each subvector independently, it follows that

$$p(\mathbf{x}'_j|\omega_i, \theta) = p(\mathbf{d}'_j|\omega_i, \theta_d)p(\mathbf{u}'_j|\omega_i, \theta_u), \quad (4)$$

where θ_d, θ_u are the network/user distortion models, respectively. Each of them contains multiple PMFs (probability mass functions) that we elaborate on below. Since in this section we consider only the network component, we assume that $p(\mathbf{u}'_j|\omega_i, \theta_u) = 1$ for all i, j , i.e., all observed user features are the same and thus perfectly match all fingerprints.

To understand the notation involved in expanding the first factor in (4), examine Fig. 3 where a host with network signature \mathbf{d}_i generates an observation \mathbf{d}'_j . Measurement begins with a SYN packet, which takes some time to get to the target, followed by the server “think” delay before it generates the

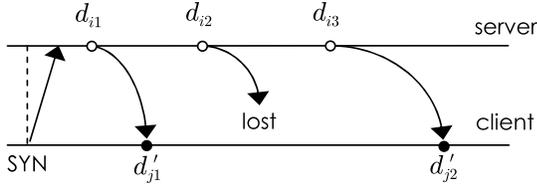


Fig. 3. Delay features (stack ω_i produces observation x'_j).

first SYN-ACK response. Database feature vectors \mathbf{d}_i consist of departure timestamps from the server, where $d_{i1} = 0$ for all i . Note that $d_{i,r+1} - d_{i,r}$ is the r -th retransmission timeout (RTO) of the stack, which was commonly used in early estimators [6], [46], [55]. Recently, however, usage of absolute timestamps $d_{i,r}$ was identified [45] as having certain modeling advantages, which is our approach as well.

On the client side, arrival timestamps $d'_{j,r}$ are measured relative to the transmission time of the SYN. Assume T_j represents the sum of the forward delay, server think time, and propagation/transmission delays of the reverse path, where T_j has some unknown distribution $f_T(\tau) = P(T_j = \tau)$. Furthermore, let $\Delta_{j1}, \Delta_{j2}, \dots$ be iid queuing delays of the return path, with another unknown distribution $f_\Delta(\delta) = P(\Delta_{j,r} = \delta)$. Then, assuming no loss, $d'_{j,r} = T_j + d_{i,r} + \Delta_{j,r}$.

To handle packet loss, assume that γ_j is a random vector that maps the received packets in observation j to their order on the server, i.e., $\gamma_j(r) = k$ means that the r -th received packet was originally in position k . In Fig. 3, for example, we have $\gamma_j = (1, 3)$. Then, if the j -th observation comes from system ω_i , it follows that

$$d'_{j,r} = T_j + d_{i,\gamma_j(r)} + \Delta_{j,r}, \quad r = 1, 2, \dots, |\mathbf{d}'_j|. \quad (5)$$

As in prior work [6], [45], [46], we assume no reordering due to the large spacing between the packets (often several seconds), which implies $\gamma_j(r+1) > \gamma_j(r)$. Let $\Gamma(i, j)$ be the set of all monotonic loss vectors that start with $|\mathbf{d}_i|$ packets and finish with $|\mathbf{d}'_j|$. Then, the Hershel+ network classifier uses $p(\mathbf{d}'_j | \omega_i, \theta_d)$ equal to [45]

$$\sum_{\tau} f_T(\tau) \sum_{\gamma \in \Gamma(i, j)} p_i(\gamma) \prod_{r=1}^{|\mathbf{d}'_j|} f_\Delta(d'_{j,r} - \tau - d_{i,\gamma(r)}), \quad (6)$$

where $p_i(\gamma)$ is the probability to observe loss pattern γ under $|\mathbf{d}_i|$ transmitted packets. For lack of a better assumption, Hershel+ uses binomial $p_i(\gamma)$, Erlang(2) $f_T(\tau)$, and exponential $f_\Delta(\delta)$, all with some fixed parameters. Since θ_d encapsulates the set of these distributions, our next goal is to recover them using EM iteration.

B. Intuition

We start with a heuristic explanation of the proposed update formulas, which is followed by a more rigorous treatment. Recall that $f_T^t(\tau)$ is an estimate of $P(T_j = \tau)$ during iteration t . Then, one obvious approach is to set this value to the average probability that each observation j has experienced a forward latency τ , conditioned on the previous estimates of unknown

parameters, i.e.,

$$f_T^{t+1}(\tau) = \frac{1}{m} \sum_{j=1}^m P(T_j = \tau | \mathbf{d}'_j, \theta_d^t, \alpha^t). \quad (7)$$

Next, each database signature with k original packets admits $2^k - 1$ unique loss patterns γ , where k goes as high as $k_{max} = 21$ in the most recent effort in the field [45]. Estimating the probability $p_i(\gamma)$ for each possible option γ is likely to produce too many unknown variables and lead to poor convergence of EM. Instead, suppose that all $\binom{k}{\ell}$ patterns of losing ℓ packets out of k are equally likely and define the probability of this event to be $q_k(\ell)$, where $k = 1, 2, \dots, k_{max}$. The resulting reduction in the number of unknown variables is significant – from roughly $2^{k_{max}+1} = 4M$ to just $k_{max}(k_{max}-1)/2 = 210$. Despite its simplicity, the framework of using $q_k(\ell)$ allows more general scenarios than the traditional iid Bernoulli model used in previous literature [45], [46].

To update distribution $q_k(\ell)$, our approach involves computing the probability that observations experienced loss of ℓ packets out of k transmitted, normalized by the probability that the original host sent k packets in the first place. To express this mathematically, define Y_j to be the number of SYN-ACKs originated by the host in observation j . Letting $\mathbf{1}_X$ be an indicator of event X , we get

$$q_k^{t+1}(\ell) = \frac{\sum_{j=1}^m P(Y_j = k | \theta_d^t, \alpha^t) \mathbf{1}_{|\mathbf{d}'_j| = k - \ell}}{\sum_{j=1}^m P(Y_j = k | \theta_d^t, \alpha^t) \mathbf{1}_{|\mathbf{d}'_j| \leq k}}, \quad (8)$$

from which the estimated probability of pattern γ is given by

$$p_i^t(\gamma) = \frac{q_{|\mathbf{d}_i|}^t(|\mathbf{d}_i| - |\gamma|)}{\binom{|\mathbf{d}_i|}{|\gamma|}}. \quad (9)$$

Finally, updates to PMF $f_\Delta^t(\delta)$ involve computing the probability that one-way delay of each received packet equals δ , normalized by the total number of packets collected during the scan, i.e.,

$$f_\Delta^{t+1}(\delta) = \frac{\sum_{j=1}^m \sum_{s=1}^{|\mathbf{d}'_j|} P(\Delta_{j,s} = \delta | \mathbf{d}'_j, \theta_d^t, \alpha^t)}{\sum_{j=1}^m |\mathbf{d}'_j|}. \quad (10)$$

C. Analysis

To make the framework outlined above usable, our next task is to express the probability of events that cannot be directly observed (e.g., $Y_j = k$, $\Delta_{j,r} = \delta$) using a recurrence that depends on only the distributions contained in θ_d^t , i.e., $(f_T^t, f_\Delta^t, q_k^t)$. Let

$$\delta_{ij\tau\gamma r} = d'_{j,r} - \tau - d_{i,\gamma(r)} \quad (11)$$

be the one-way delay $\Delta_{j,r}$ conditioned on $T_j = \tau$, loss pattern γ , signature ω_i , and observation j . For brevity of notation, suppose $\sum_{ij\tau\gamma s}$ refers to five nested summations, where i goes from 1 to n , j rolls from 1 to m , τ moves over all bins of the PMF $f_T(\tau)$, γ iterates over all monotonic loss vectors in $\Gamma(i, j)$, and s travels from 1 to $|\mathbf{d}'_j|$. If some of the variables are absent from the subscript,

TABLE IV
CLASSIFICATION RESULTS OF NETWORK EM IN D_1

Case	ρ^1	ρ^∞	α^∞
S_{11}	0.67	0.95	0.90, 0.05, 0.05
S_{12}	0.48	0.91	0.05, 0.90, 0.05
S_{13}	0.45	0.95	0.90, 0.05, 0.05
S_{14}	0.60	0.85	0.30, 0.60, 0.10

the corresponding sums are omitted from the result. With this in mind, define

$$p_{ij\tau\gamma}^t := \alpha_i^t f_T^t(\tau) p_i^t(\gamma) \prod_{r=1}^{|\mathbf{d}'_j|} f_\Delta^t(\delta_{ij\tau\gamma r}), \quad (12)$$

$$\beta_{ij\tau\gamma}^t := p(\omega_i, \tau, \gamma | \mathbf{d}'_j, \theta_d^t, \alpha^t) = \frac{p_{ij\tau\gamma}^t}{\sum_{i\tau\gamma} p_{ij\tau\gamma}^t} \quad (13)$$

and consider the next result.

Theorem 2: Under network distortion, estimators (3), (7), (8), and (10) can be written as

$$\alpha_i^{t+1} = \frac{1}{m} \sum_{j\tau\gamma} \beta_{ij\tau\gamma}^t, \quad (14)$$

$$f_T^{t+1}(\tau) = \frac{1}{m} \sum_{ij\gamma} \beta_{ij\tau\gamma}^t, \quad (15)$$

$$q_k^t(\ell) = \frac{\sum_{ij\tau\gamma} \beta_{ij\tau\gamma}^t \mathbf{1}_{|\mathbf{d}'_j|=k-\ell, |\mathbf{d}_i|=k}}{\sum_{ij\tau\gamma} \beta_{ij\tau\gamma}^t \mathbf{1}_{|\mathbf{d}'_j| \leq |\mathbf{d}_i|=k}}, \quad (16)$$

$$f_\Delta^t(\delta) = \frac{\sum_{ij\tau\gamma s} \beta_{ij\tau\gamma}^t \mathbf{1}_{\delta_{ij\tau\gamma s}=\delta}}{\sum_j |\mathbf{d}'_j|}. \quad (17)$$

While the result of Theorem 2 may appear daunting due to the number of nested summations, its implementation in practice can be done with little extra cost compared to Hershel+. Specifically, usage of (6) in (1) for all i, j already requires five nested loops. In the inner-most loop of that algorithm, (17) adds one increment to a hash table that maintains the PMF of one-way delay. Updates in (14)-(16) are performed less frequently and, in comparison, consume negligible time. The only caveat is that Hershel+ can be optimized [45] to remove the outer summation in (6) when f_T is Erlang(2) and f_Δ is exponential. Our approach, on the other hand, requires a hash-table lookup for both distributions. This makes its single iteration similar in speed to unoptimized Hershel+.

Theorem 3: Iteration (14)-(17) is the EM algorithm for (θ_d, α) .

D. Discussion

We revisit earlier simulations on dataset D_1 , run (14)-(17) over the same input, and show the result in Table IV. Compared to Table III, the derived EM estimator significantly improves the accuracy of both classification and vector α . Note that S_{12} contain 43% of the observations with just one packet, i.e., zero RTOs. In methods that rely on RTO [6], [46], [55], these samples would be either discarded as impossible to classify or assigned to a uniformly random signature. In contrast, estimator (14)-(17) manages to do much better as it learns distributions (f_T, f_Δ, α) and makes the best decision possible under these conditions. The accuracy of estimated

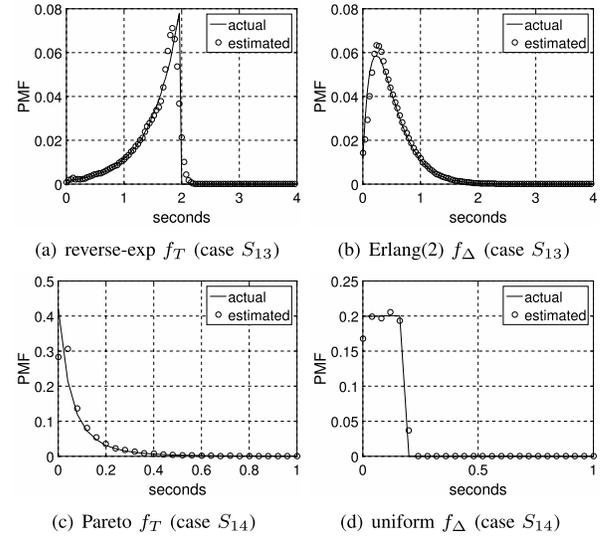


Fig. 4. Recovery of delay parameters in D_1 .

delay distributions is shown in Fig. 4. With the exception of noise at the points of discontinuity of each density, functions $f_T^\infty, f_\Delta^\infty$ match the true parameters quite well.

Recalling (5), where $T_j + \Delta_{jr}$ are always measured together, it may not be obvious how T_j can be separated from Δ_{jr} and why the result in Fig. 4 is possible. Indeed, this is reminiscent of the classical deconvolution problem: given observations $\{X_i + Y_i\}_{i=1}^m$, where $X_i \sim F_X(x)$ and $Y_i \sim F_Y(x)$ are iid, determine the individual distributions F_X, F_Y . Deconvolution is generally unsolvable unless either F_X or F_Y is known ahead of time. While our problem is similar, there is a crucial difference – EM can see the same value T_j coupled with *multiple* instances of Δ_{jr} , for $r = 1, 2, \dots, |\mathbf{d}'_j|$. As long as $q_k(k-1) < 1$ (i.e., packet loss leaves at least two packets in enough observations) and $m \rightarrow \infty$, deconvolution is possible in our setting, but up to a location shift, i.e., one of the estimated distributions may be shifted left by a constant and the other right by the same amount. If we know that one of them starts at zero, it is possible to determine the shift after the fact. Furthermore, if both estimated densities $f_T^\infty, f_\Delta^\infty$ already begin at zero, no correction is needed. This is the case in Fig. 4 and later in our Internet scan.

Since all signatures in D_1 had three packets, it was easy to figure out the number of them lost in each \mathbf{d}'_j , which led to q_k^∞ being perfectly accurate, regardless of whether (16) was used or not. In a more interesting database, which we call D_2 , Linux is augmented with a fourth packet that follows after a 3-second RTO. To experiment with different loss patterns, define $\text{BinT}(k, p)$ to be a binomial distribution with parameters (k, p) truncated to the range $[0, k-1]$. Since the loss of all k packets cannot be observed, we avoid generating this case in the simulator. Additionally, suppose $\text{RevBin}(k, p)$ is the *reverse binomial distribution* such that $X \sim \text{BinT}(k, p)$ and $Y = k-1-X$ implies $Y \sim \text{RevBin}(k, p)$. This is used in scenario S_2 from Table V, which shows q_k and the average observed loss rate among the signatures with k packets.

Table VI shows classification results for three methods – Hershel+, the partial EM framework without loss updates (16), and the full algorithm from Theorem 2. Not surprisingly,

TABLE V
NETWORK PARAMETERS OF SCENARIO S_2

Case	Delay	q_3	Loss	q_4	Loss
S_{21}	As in S_{12}	BinT(3, 0.3)	28%	BinT(4, 0.3)	30%
S_{22}	As in S_{12}	BinT(3, 0.1)	10%	BinT(4, 0.8)	66%
S_{23}	As in S_{12}	RevBin(3, 0.1)	57%	RevBin(4, 0.1)	65%
S_{24}	As in S_{13}	BinT(3, 0.7)	54%	BinT(4, 0.7)	61%

TABLE VI
CLASSIFICATION RESULTS IN D_2

Case	α	Hershel+		EM α, f_T, f_Δ		Full EM	
		ρ^1	α^1	ρ^∞	α^∞	ρ^∞	α^∞
S_{21}	0.90	0.76	0.68	0.70	0.63	0.91	0.90
	0.05		0.25		0.31		0.05
	0.05		0.07		0.05		0.05
S_{22}	0.90	0.45	0.34	0.13	0.06	0.97	0.90
	0.05		0.47		0.84		0.05
	0.05		0.19		0.10		0.05
S_{23}	0.90	0.45	0.36	0.10	0.06	0.90	0.90
	0.05		0.46		0.90		0.05
	0.05		0.18		0.04		0.05
S_{24}	0.90	0.42	0.33	0.14	0.10	0.92	0.90
	0.05		0.38		0.88		0.05
	0.05		0.29		0.02		0.05

Hershel+ again struggles to recover α , even when its classification accuracy is pretty high. Omission of (16) does create challenges for partial EM, where in all four cases it produces worse results than Hershel+. On the other hand, the full algorithm improves accuracy and delivers the exact α despite complex underlying network conditions.

V. USER FEATURES

A. Distortion Model

Our goal in this section is to expand the second factor in (4) and develop an estimator for its distortion model. This is done in isolation from the network features, i.e., using $p(\mathbf{d}'_j|\omega_i, \theta_d) = 1$ for all i, j . Assume $b \geq 1$ user features, where each observation j provides a constant-length vector $\mathbf{u}'_j = (u'_{j1}, \dots, u'_{jb})$. These include the TCP window size, IP TTL (time to live), IP DF (do not fragment), TCP MSS (maximum segment size), and TCP options, for a total of $b = 5$ integer-valued fields. Since RST features depend on network loss, we delay their discussion until the next section. Note that each field may be allocated a different number of bits in the header and the number of available options a_v for u'_{jv} may depend on v (e.g., two for DF and 64K for Win).

Modification to user features at the target host, which we model with a set of distributions θ_u , can be accomplished by manually changing default OS parameters (e.g., editing the registry), using specialized performance-tuning software, requesting larger/smaller receiver kernel buffers while waiting on sockets (i.e., using `setsockopt`), and deploying network/host scrubbers [12], [40], [42], [51], [56] whose purpose is to obfuscate the OS of protected machines. Besides intentional feature modification, distortion θ_u may also accommodate unknown network stacks that build upon a documented OS, but change some of its features (e.g., new versions of embedded Linux customized to a particular device).

Prior work either omits formally modeling user volatility [6], [7], [39], [58], [60] or assumes that u_{iv} stays the same

with some probability π_v and changes to another integer with probability $1 - \pi_v$ [45], [46]. While the latter approach works well in certain cases, it has limitations. Besides the fact that π_v is generally unknown, binary decision-making fails to create a distribution over the available choices. For example, $\pi_v = 0.9$ assumes that *each* of the 65,534 non-default window sizes occurs with probability 0.1. Instead, a more balanced approach is to normalize $1 - \pi_v$ by the number of available distortion options $a_v - 1$. Second, it is likely that certain devices are modified less frequently than others (e.g., due to firmware restrictions) and individual distortions are OS-specific, which implies that π_v should depend on i . Finally, the existing methods have no way of tracking the location and probability mass of distortion, which does not have to be uniform in practice (e.g., window size 257 bytes is less likely than 64K).

To overcome these problems, assume that $\pi_{iv}(y)$ is the probability that feature v of OS i is modified to become y , which gives rise to a set of nb distributions that comprise our user-distortion model θ_u . Then, the proposed classifier is

$$p(\mathbf{u}'_j|\omega_i, \theta_u) = \prod_{v=1}^b \pi_{iv}(u'_{jv}), \quad (18)$$

where modification to features is assumed to be independent. Note that doing otherwise does not appear tractable at this point (i.e., estimation of covariance matrices produces too many variables for EM to handle).

B. Iteration

We begin by discussing under what conditions the problem is identifiable, despite having a large number of unknown distributions. Assume $\phi_{iv} := \pi_{iv}(u_{iv})$ is the probability with which feature v stays the same for OS i . Because we do not know ahead of time the reasoning of the user for changing the features or the new values of modified fields, the estimation problem for π_{iv} is unsolvable unless enough of the probability mass remains at the original location, i.e., ϕ_{iv} is above some threshold. From common sense, it is likely that $\phi_{iv} \geq 0.5$ holds among the general population of Internet hosts; however, EM converges under even weaker conditions (e.g., when ϕ_{iv} is the largest value in each PMF π_{iv}). Coupling this with an initial state that satisfies the same constraint allows EM to discover a unique solution.

We define the estimator for user distortion as the probability to observe y in feature v across all matches against OS i , i.e.,

$$\pi_{iv}^{t+1}(y) = \frac{\sum_{j=1}^m p(\omega_i|\mathbf{u}'_j, \theta_u^t, \alpha^t) \mathbf{1}_{u'_{jv}=y}}{\sum_{j=1}^m p(\omega_i|\mathbf{u}'_j, \theta_u^t, \alpha^t)}. \quad (19)$$

To allow simplification of this expression below, define

$$p_{ij}^t := \alpha_i^t p(\mathbf{u}'_j|\omega_i, \theta_u^t, \alpha^t) = \alpha_i^t \prod_{v=1}^b \pi_{iv}^t(u'_{jv}), \quad (20)$$

$$\beta_{ij}^t := p(\omega_i|\mathbf{u}'_j, \theta_u^t, \alpha^t) = \frac{p_{ij}^t}{\sum_{i=1}^n p_{ij}^t}. \quad (21)$$

The next result follows from substitution of (20)-(21) into (3) and (19), as well as earlier proofs of Theorems 1 and 2.

TABLE VII
USER FEATURES OF DATABASE D_3

OS	Win	TTL	DF	OPT	MSS
Linux	5,792	64	1	MSTNW	1,460
Windows	16,384	128	0	MNWNNTNNS	1,380
Novell	6,144	128	1	MNWSNN	1,460

Theorem 4: Under user distortion, estimators (3) and (19) can be written as

$$\alpha_i^{t+1} = \frac{1}{m} \sum_{j=1}^m \beta_{ij}^t, \quad (22)$$

$$\pi_{iv}^{t+1}(y) = \frac{\sum_{j=1}^m \beta_{ij}^t \mathbf{1}_{u'_{jv}=y}}{m\alpha_i^{t+1}}. \quad (23)$$

Furthermore, this is the EM algorithm for (θ_u, α) .

C. Discussion

To evaluate the result of Theorem 4, we construct a new database D_3 , shown in Table VII, by switching from RTOs to user features. Note that this Linux signature ties Novell in DF and MSS, while Windows does the same in TTL. For simplicity of presentation, we use scenarios with $\phi_{iv} = \phi_v$ for all i , where ϕ_v is the probability with which feature v stays at the default value. This replaces matrix ϕ_{iv} with a vector ϕ_v , which is easier to follow across the different tables.

The initial PMFs π_{iv}^0 of EM are set up to include 90% of the mass on the default value and split the remainder uniformly across the viable alternatives. Since it is believed [46] that the order of non-NOP options cannot be changed without rewriting the TCP/IP stack of the OS, we initialize π_{i4}^0 to allow only candidates compatible with the original u_{i4} . For example, MST is feasible for Linux, but not the other two signatures in Table VII. Note that any single option (M, S, W) and the empty set are valid for all three OSes.

We use two models for generating noisy observations. The first one, which we call RAND, picks uniformly from the space of possible values observed in our Internet scan, except OPT is limited to compatible subsets/supersets of the original. We have 5,695 candidates for Win, four for TTL, two for DF, 266 for OPT, and 1,082 for MSS. Decisions are made independently for each feature v and each observation j , which models users “tweaking” their OS without coordinating with each other or sharing a common objective. Although RAND can generate 13.1 billion unique combinations \mathbf{u}'_j , only a small subset is encountered by the classifier in simulations below.

The second model, which we call PATCH, selects an alternative vector of features \mathbf{u}''_i for each OS ω_i and switches the default value u_{iv} to u''_{iv} with probability $1 - \phi_v$, again independently for each v . This represents deployment of software patches that change one of the features to an updated value. The probability for a host to use multiple patches is the product of corresponding $(1 - \phi_v)$'s. For example, modification to both Win and OPT affects $(1 - \phi_1)(1 - \phi_4)$ fraction of hosts. Vectors \mathbf{u}''_i are non-adversarial and do not attempt to confuse the classifier. We construct them by flipping the DF flag, setting OPT to M, and adding i to all remaining fields (modulo the max field value). The result is given in Table VIII.

TABLE VIII
PATCHED USER FEATURES

Vector	Win	TTL	DF	OPT	MSS
\mathbf{u}''_1	5,793	128	0	M	1,461
\mathbf{u}''_2	16,386	32	1	M	1,382
\mathbf{u}''_3	6,147	64	0	M	1,463

TABLE IX
PARAMETERS OF SCENARIO S_3

Case	Model	Feature stay prob ϕ_v	Popularity α
S_{31}	RAND	(0.3, 0.2, 0.5, 0.4, 0.4)	(0.90, 0.05, 0.05)
S_{32}	RAND	(0.0, 0.0, 0.1, 0.2, 0.0)	(0.90, 0.05, 0.05)
S_{33}	PATCH	(0.2, 0.2, 0.2, 0.2, 0.2)	(0.7, 0.2, 0.1)

TABLE X
CLASSIFICATION RESULTS IN D_3

Case	Hershel+	EM			
	ρ^1	ρ^1	ρ^∞	ϕ_v^∞	
S_{31}	0.76	0.79	0.96	(0.30, 0.20, 0.50, 0.40, 0.40)	
S_{32}	0.29	0.32	0.91	(0.00, 0.00, 0.10, 0.20, 0.00)	
S_{33}	0.31	0.50	1	(0.20, 0.20, 0.20, 0.20, 0.20)	

To estimate vector ϕ_v^t in the classifier, we use a weighted average of feature non-modification across all OSes, i.e., $\phi_v^t = \sum_{i=1}^n \alpha_i^t \phi_{iv}^t$. Our next scenario S_3 is detailed in Table IX and the corresponding outcome is given by Table X. We omit vector α^∞ since it matches ground-truth α very accurately. Due to the new treatment of non-default features in (18), the first iteration of EM in Table X is superior to Hershel+. However, both are much worse than the last iteration. It should be noted that the second case S_{32} modifies Win, TTL, and MSS in 100% of the samples. Identifiability in such conditions is helped by the fact that OPT is constrained to a subset of the original string, which makes a certain fraction of randomly generated values feasible for only one OS. This allows EM to learn to ignore (Win, TTL, MSS) and focus decisions on (DF, OPT). Furthermore, when guessing is involved, EM uses its knowledge of α to correctly pick the most-likely OS. It is also interesting that S_{33} is classified with 100% accuracy once EM gets a grasp on the new values in Table VIII and their probability of occurrence.

VI. COMPLETE SYSTEM

A. Reset Packets

Because loss of RST packets causes the corresponding user features (i.e., ACK/RST flags, ACK sequence number, window size [46]) to be wiped out, there is dependency between distortion applied by the network and the user. As a result, this case should be handled separately. The first modification needed is to increase the length of network vectors \mathbf{d}_i and \mathbf{d}'_j to accommodate the RST timestamp. The second change is to add RST values into user features. Since RST fields are often implementation choices/bugs that cannot be modified independently of each other [46], they can be combined into a single integer and appended to user vectors \mathbf{u}_i and \mathbf{u}'_j in position $b + 1$.

TABLE XI
HANDLING OF RST PACKETS

RST present		Action	Multiplier ζ_{ij}^t
\mathbf{d}'_j	\mathbf{d}_i		
yes	yes	–	$\pi_{i,b+1}^t(u'_{j,b+1})$
yes	no	ignore RST in \mathbf{d}'_j	$\pi_{i,b+1}^t(u'_{j,b+1})$
no	yes	–	1
no	no	–	1

There are four possible scenarios for handling RST packets. They are shown in Table XI, each with a certain probability ζ_{ij}^t that must be factored into the formulas developed earlier. When both the observation and candidate signature contain a RST, the only multiplier needed is the probability that the received feature was produced by that OS. If the sampled OS has a RST, but the signature does not, this indicates a possible interference from an intermediate device (e.g., IDS after expiring connection state, scrubbers). In this case, it is likely meaningless to use the temporal characteristics of the RST, which is why we omit it from \mathbf{d}'_j before computing the loss and delay probabilities. However, multiplication by $\pi_{i,b+1}^t(u_{j,b+1})$ is still warranted since we must assign a proper weight to this mismatch. The third row of the table corresponds to packet loss, which is handled automatically in $p_i^t(\gamma)$, i.e., no additional actions or multipliers are needed. Finally, the last row is identical to the setup assumed in preceding sections.

B. Final Model

We now combine the developed network, user, and RST models into a single framework. Redefining (12) as

$$p_{ij\tau\gamma}^t = \alpha_i^t \zeta_{ij}^t \left(\prod_{v=1}^b \pi_{iv}^t(u'_{jv}) \right) f_T^t(\tau) p_i^t(\gamma) \prod_{r=1}^{|\mathbf{d}'_j|} f_{\Delta}^t(\delta_{ij\tau\gamma r}) \quad (24)$$

allows us to compute $\beta_{ij\tau\gamma}^t$ still via (13), as well as reuse (14)-(17). However, (23) requires an update to

$$\pi_{iv}^{t+1}(y) = \frac{\sum_{j=1}^m \mathbf{1}_{u_{jv}=y} \sum_{\tau\gamma} \beta_{ij\tau\gamma}^t}{m \alpha_i^{t+1}}, \quad (25)$$

where $v = 1, 2, \dots, b+1$. The final classifier, which we call *Faulds*,¹ is applied after EM has converged and is given by

$$p(\omega_i | \mathbf{x}'_j, \theta^\infty, \alpha^\infty) = \sum_{\tau\gamma} \beta_{ij\tau\gamma}^\infty. \quad (26)$$

It is easy to generalize our earlier results to cover the complete model, as given in the next theorem without proof.

Theorem 5: Under both network and user distortion, estimator (13)-(17), (24)-(25) is the EM algorithm for (θ, α) .

C. Scaling the Database

Due to the large number of features it combines, *Faulds* is not challenged by the previous toy databases. We therefore switch to a more realistic set of signatures created by Plata

¹Henry Faulds was a Scottish scientist who extended the ideas of William Herschel and proposed the first usable forensic fingerprint-identification method in 1880.

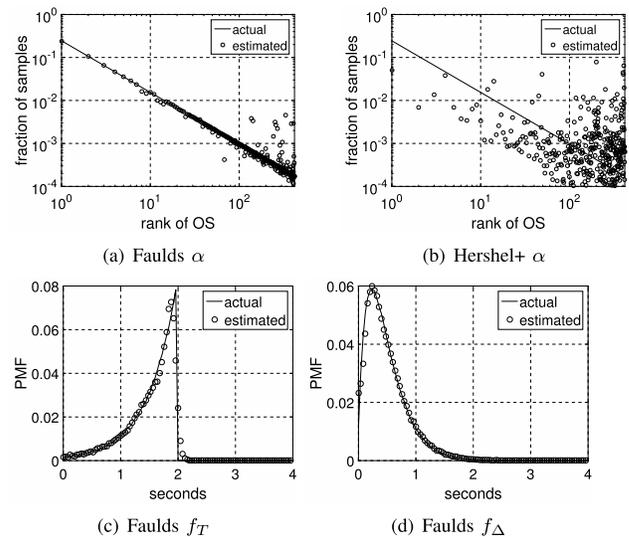


Fig. 5. Results in D_4 .

in [45]. We call this database D_4 and note that it contains 420 stacks, among which some have the same exact RTO vector and others overlap in *all* user features. The database was constructed to ensure that signatures were sufficiently unique under delay distortion, but packet loss and user modifications were not taken into account. As a result, the database contains a number of entries that would be difficult to distinguish under the types of heavy distortion considered in this paper. Nevertheless, these tests should indicate how well *Faulds* scales to larger databases and whether its recovery of the unknown parameters (α, θ) is affected by an increased uncertainty during the match.

We set popularity α to the Zipf distribution with shape parameter 1.2 and continue using $m = 2^{18}$ observations, which gives us 64K samples from the most common OS and just 49 from the least. We borrow the delay from case S_{13} (i.e., reverse-exponential T with mean 1.5 sec, Erlang(2) Δ with mean 0.5) and packet loss from S_{23} (i.e., reverse-binomial). Finally, we use RAND with stay probability $\phi_v = 0.8$.

The first iteration of *Faulds* produces a respectable $\rho^1 = 0.42$. This is gradually improved with each step, until convergence to a more impressive $\rho^\infty = 0.70$. To make sense out of α^∞ , we sort all signatures in rank order from the most popular to the least and plot the result in Fig. 5(a). There is a strong match in the top-100, while the random noise in the tail is explained by the scarcity of these OSes in the observation (i.e., below 250 samples each). For comparison, the outcome of Hershel+ is displayed in part (b) of the figure. To complete the big picture, subfigures (c)-(d) show estimates of f_T and f_{Δ} . Despite an overall 30% classification mismatch, these PMFs are no worse than previously observed in Fig. 4, which indicates that incorrect decisions overwhelmingly went to signatures with similar RTO vectors as the true OS.

Instead of scrutinizing 21 different loss PMFs, suppose we compute a single metric – the fraction of packets dropped within the entire observation \mathbf{x}' , conditioned on at least one

TABLE XII
INJECTION CLASSIFICATION SUMMARY

Size of D_4	Injected samples	ρ_*^t	ρ_*^∞	p_{loss}^∞	$E[\phi_v^\infty]$
378 (90%)	7,089 (2.8%)	0.88	0.91	0.10	0.80
336 (80%)	49,648 (19.0%)	0.87	0.89	0.11	0.74
294 (70%)	60,058 (22.9%)	0.87	0.89	0.11	0.73
210 (50%)	91,408 (34.9%)	0.91	0.91	0.11	0.72
126 (30%)	189,293 (72.2%)	0.95	0.93	0.17	0.60

packet surviving. To this end, define during step t

$$L_k^t = \sum_{\ell=1}^{k-1} \ell q_k^t(\ell) \quad (27)$$

to be the average number of lost replies in signatures with k packets. Then, taking an estimated ratio of all dropped packets to the total transmitted yields the expected loss rate

$$p_{loss}^t = \frac{\sum_{i=1}^n \alpha_i^t L_{|d_i|}^t}{\sum_{i=1}^n \alpha_i^t |d_i|}. \quad (28)$$

Recall that the simulation allowed loss to affect at most $k-1$ packets in OSes with $|d_i| = k$. Therefore, its ground-truth packet loss should represent the same quantity as (28). Traces show that 70.1% of the packets were dropped, which matches quite well against $p_{loss}^\infty = 69.3\%$.

Since $\phi_v = 0.8$ was a constant in this simulation, it makes sense to compare it against feature-modification estimates averaged across all fields and all OSes, i.e.,

$$E[\phi_v^t] = \frac{1}{b+1} \sum_{v=1}^{b+1} \phi_v^t = \frac{1}{b+1} \sum_{v=1}^{b+1} \sum_{i=1}^n \alpha_i \phi_{iv}^t. \quad (29)$$

Results show that $E[\phi_v^\infty] = 0.802$, which is very close to the actual value. While there is some variation in individual ϕ_{iv} , it is of little concern due to the small number of samples seen by Faulds from these OSes.

D. Unknown Signatures

We recognize that having a database that knows all devices on the Internet is near impossible. Therefore, infiltration of samples from unknown signatures into \mathbf{x}' , which we call *injections*, is inevitable in practice. Understanding their impact is our next topic.

Suppose \mathbf{x}'_j is produced by some unknown OS ω that does not belong to the database. If \mathbf{x}'_j is so different from the known signatures that $p(\mathbf{x}'_j | \theta^t, \alpha^t) = 0$, i.e., it matches each OS with probability 0, its injection into the observation will contribute nothing to updates of (α^t, θ^t) and thus will have no impact on classification decisions. In order to achieve a flat-out mismatch of this type, either delay $\delta_{ij\tau\gamma}$ must be negative for all i, τ, γ or the product in (24) must be smaller than the precision of floating-point arithmetic.

For injections with $p(\mathbf{x}'_j | \theta^t, \alpha^t) > 0$ the situation is less clear-cut. For such cases, we want the classifier to select the statistically most likely OS among the database, which Faulds naturally does after EM convergence. Almost all injected cases in this section fall into this category; however, because there is no easy way to evaluate such matches, we do not focus on

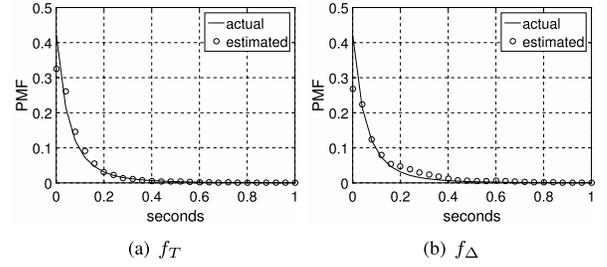


Fig. 6. Recovered delay under 72% injection.

them here. Instead, we are interested in classification accuracy of the samples generated by the *known* OSes, where the result may become affected by the EM feedback loop that learns potentially incorrect information from injected samples.

When the unknown OS ω is close to an existing signature ω_i , injections are minimally different from distorted instances of \mathbf{x}_i . As a result, they do not negatively impact EM or its convergence point. On the other hand, it is also possible that \mathbf{x}'_j is a potential match to multiple unrelated OSes and the amount of distortion needed to make them appear as \mathbf{x}'_j is much greater than the underlying θ . If the volume of injections is high, how likely is EM to introduce bias into distributions of delay/loss to the point of impacting classification accuracy for *non-injected* samples?

We do not consider encountering of adversarial injections (i.e., special signatures crafted to cause maximum harm for a given database and classifier) to be likely in practice and instead focus on evaluating the effect of random subset removal from D_4 . Specifically, assume the simulator produces distorted observations using all 420 network stacks; however, Faulds has access to only some of the original signatures. For the next simulation, we use Pareto f_T and f_Δ , both with mean 0.1 seconds, iid packet loss at 10%, and $\phi_v = 0.8$.

Define ρ_*^t to be the classification accuracy among non-injected observations during step t and consider Table XII, which shows the shrunk database size, number of injected samples among $m = 2^{18}$ observations, and the output of Faulds. The result shows that removal of signatures does not carry a significant negative impact on accuracy of classification for the known OSes. In fact, ρ_*^t slightly rises as the database shrinks since it becomes easier to classify among fewer options. Packet loss p_{loss}^∞ also appears immune, except in the last row where 72% of \mathbf{x}' contains observations from unknown OSes. Its increase to 17% is explained by more frequent matches that require high packet loss to be feasible. Finally, the feature-stay probability in the last column is the most affected, which was also expected due to the increased header-field mismatch.

Fig. 6 shows the two delay PMFs estimated by Faulds in the last row of Table XII. Recovery is quite good, except for a slight bump in f_Δ between 200 and 400 ms. This shows that removing 70% of the signatures in D_4 still leaves enough unique RTO vectors to produce highly accurate results. In the actual Internet, however, we do not expect injection conditions to be anywhere near these levels because D_4 contains an array of major network stacks (e.g., Windows, Unix), printer

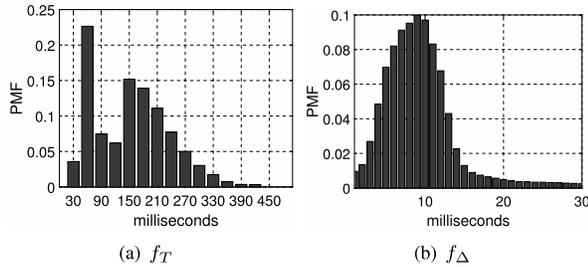


Fig. 7. Internet delay distributions.

firmware (e.g., HP, Lexmark, Brother), Cisco equipment, and various derivative implementations that run on embedded devices. See [45] for more details.

VII. INTERNET MEASUREMENT

This section illustrates how Faults can be applied in practice and the types of information it produces. Since the algorithms proposed earlier in the paper are based on fundamental properties of fingerprints and rigorous statistical theory, without being tied to any particular training input or a-priori assumptions, Faults is guaranteed to converge to the optimal (in the EM sense) result given the available information. This applies to measurements collected in the past, present, and future, including wide-area networks, LANs, and other scenarios. As a result, it is sufficient to discuss any one Internet snapshot and leave other cases for the future, where our release of Faults source code and its database [21] should enable the community to perform additional measurements and network characterization as needed.

A. Overview

Our dataset comes from [47], which is the most recent Internet-wide scan that captured all of the parameters needed for Faults. On December 14, 2016, this experiment probed all 2.8B BGP-reachable IPv4 addresses on port 80 and gathered responses from 67.6M hosts. Using a dual 8-core AMD Opteron 6320 @ 2.8 GHz, Faults executed at 3,801 hosts per second and completed a full pass over the data in roughly 5 hours. In large-scale classification, such as the one attempted here, Faults produces a huge volume of information in the form of various PMFs and estimates. Due to limited space, we present only a brief review of the obtained results and leave more detailed analysis (including attempts to uncover injections and correct for them) for future work. We start with basic sanity checks of the estimated distortion θ and then delve into classification result α .

B. Network Distortion

Fig. 7(a) shows the recovered distribution f_T using bin size 30 ms. Delays below 60 ms (29%) represent unloaded servers in close proximity to the scanner, most likely within the continental US. Those in the 120 – 200 ms range (40%) indicate targets whose RTTs are consistent with destinations in Europe and Asia. The remaining cases covers longer paths, OS scheduling delays, non-trivial CPU load on the server,

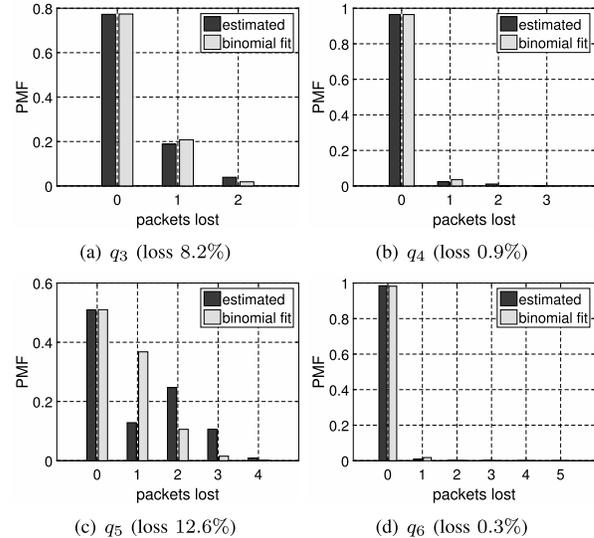


Fig. 8. Internet packet-loss PMFs.

and involvement of various backend databases to set up the connection. Overall, we obtain $E[T_j] = 148$ ms, 80% of the samples below 200 ms, and 99.2% below 450 ms. Fig. 7(b) plots the distribution of one-way delay f_Δ , in which 92% of the mass concentrates below 30 ms and 97% below 100 ms. The average queuing delay $E[\Delta_{j_r}] = 15$ ms also sounds quite reasonable.

To examine packet loss, define $\eta_k^t = \sum_{i=1}^n \alpha_i^t \mathbf{1}_{|d_i|=k}$ to be the estimated fraction of observations that use an OS with k packets. The top values of k are four ($\eta_4^\infty = 0.42$, 112 stacks in Plata database D_4), six ($\eta_6^\infty = 0.31$, 80 stacks), three ($\eta_3^\infty = 0.07$, 72 stacks), and five ($\eta_5^\infty = 0.04$, 54 stacks). Fig. 8 plots the recovered loss PMFs for these values of k , each fitted with an iid binomial model and accompanied by the average loss rate L_k^∞/k from (27). First, it is interesting that the loss rate is heterogeneous, ranging from 0.3% in q_6 to 12.6% in q_5 . This phenomenon may be inherent to the signatures that map to each k (e.g., certain printers cut the SYN-ACK sequence when their tiny SYN backlog queue overflows [45]), the load on the corresponding Oses, and host location on the Internet, all of which suggests there is an extra benefit to estimating q_k independently for different k . Second, while in a few of the plots the binomial model shows a reasonable fit, this does not universally hold. Finally, computing (28) for the Internet scan yields an average loss rate of 3.3% across all observations. This is consistent with 3.8% found in a Google study of SYN-ACK retransmission rates [11].

C. User Distortion

Computing (29), we obtain $E[\phi_v^\infty] = 0.81$, i.e., the average probability to encounter a non-default value was 19%. Faults produced $420 \times 6 = 2,520$ distributions of user features, among which we highlight several interesting cases, focusing on the two most volatile fields – Win and MSS – and limiting all PMFs to values above the 1% likelihood. Since MSS sometimes depends on the MTU of the underlying data-link

TABLE XIII
 FAULTS CLASSIFICATION AT ITERATION 1 (LEFT) AND 100 (RIGHT)

OS	α^1	Count	OS	α^{100}	Count	Change
Ubuntu / Redhat / CentOS	0.224	14,098,093	Ubuntu / Redhat / CentOS	0.334	21,361,956	0.52
Ubuntu / SUSE / CentOS	0.111	8,896,622	Embedded Linux	0.103	6,467,303	0.02
Embedded Linux	0.082	6,326,349	Windows 7 / 2008 / 2012	0.056	3,669,372	0.25
Windows 7 / 2008 / 2012	0.047	2,942,254	Schneider / APC Embedded	0.055	3,632,638	1.29
Ubuntu / Redhat / SUSE	0.037	2,408,386	Ubuntu / Redhat / SUSE	0.031	2,001,329	-0.17
Schneider / APC Embedded	0.022	1,587,396	Windows XP / 2003	0.018	1,248,619	-0.05
Windows XP / 2003	0.021	1,314,967	Redhat / CentOS / SUSE	0.016	1,046,567	-0.17
Redhat / CentOS / SUSE	0.018	1,254,797	Dell Laser / Xerox WorkCenters	0.015	976,717	0.25
Embedded Linux	0.015	1,044,028	Windows 2008 R2 / 2012	0.014	837,466	-0.08
Windows 2008 R2 / 2012	0.013	907,167	Cisco Embedded	0.013	824,039	2.29

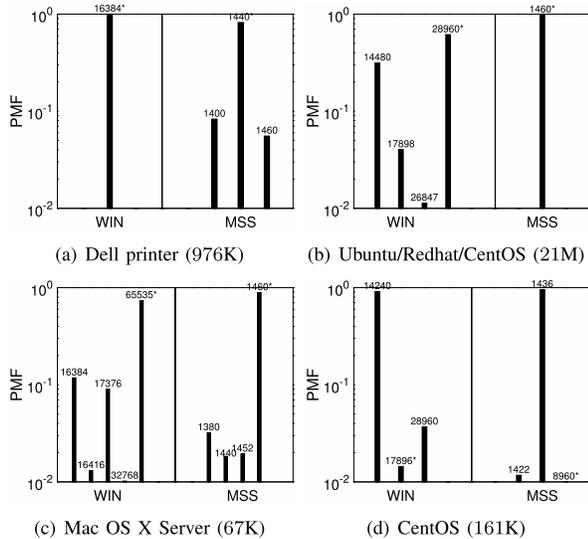


Fig. 9. Internet distributions π_{i1} and π_{i5} (default values have an asterisk).

layer and/or tunneling protocol (e.g., IPv6), this field may experience fluctuation even if the OS does not allow explicit means for changing this value.

We expected devices with firmware restrictions that prevent user access to the configuration of SYN-ACK parameters to exhibit high ϕ_{iv} . One example is shown in Fig. 9(a) for a popular Dell printer. Among 976K occurrences on the Internet, this device keeps the default window with probability 1. Intuition also suggests that general-purposes OSES are more susceptible to modification and/or existence of alternatively patched versions. One example is 21M hosts with Ubuntu Linux, where Fig. 9(b) shows that Faults discovers 31% of the cases with window size exactly half of the default (i.e., 14,480 instead of 28,960). A more dispersed case is Mac OS X Server in part (c), which exhibits noticeable variation in both Win and MSS. Its default values remain with probability 73% and 89%, respectively. Finally, in subfigure (d), CentOS (enterprise Linux) has its original combination (17,896, 8,960) occurring in only 1% of the cases. We conjecture that the Plata database [45], which was constructed from production devices in a large campus network, captured a non-standard version of this stack with jumbo Ethernet frames enabled. Since this is an inherent property of any database, it is important to allow great flexibility in the match process to accommodate such scenarios.

D. Classification Results

We define Faults to be successful for sample j if the denominator of (1) is non-zero, i.e., $p(\mathbf{x}_j|\theta^t, \alpha^t) > 0$. This means that at least one OS matches \mathbf{x}_j^t with a non-zero probability. Using the Plata database with 420 network stacks [45], Faults successfully classified 63.1M hosts (i.e., 93%). From a pure statistical point of view, the remaining 4.5M devices should be assigned to the OS with the highest α_i^∞ . But it is also likely these cases come from unknown stacks or observations with too much packet loss, in which case excluding them from classification might be prudent as well, which is our approach below.

The left side of Table XIII shows the top ten OSES after one iteration of Faults. Note that the Plata database was *auto-generated* from a pool of devices found at a university network. Even though this process [45] produced only a high-level description of each OS, additional *manual* effort can be used to provide each signature with a more specific kernel version and/or physical device. We consider this issue orthogonal to the topic of the paper since Faults operates on TCP/IP signatures and its accuracy does not depend on the name affiliated with each fingerprint \mathbf{x}_i .

The dominance of Linux and embedded devices in Table XIII (left) matches the statistics reported in prior work [27], [45], [46], although a more interesting result is the amount of relative change occurring in the classification as Faults goes through its iterations. Table XIII (right) shows the α vector after 100 steps. The top Linux signature gains 52%, Windows 7 in third place increases by 25%, and two other Linux stacks drop 17% each. Further down the list, there is significant movement as well, where certain embedded systems, such as Schneider APC (data-center hardware solutions), Dell printers, and Cisco, increase their membership by 25 – 229%. There is even more shuffle outside the top-10, which underscores the importance of using proper algorithms for estimating α .

Table XIV splits all classified hosts into eight categories. The top two signatures are desktop/server OSES and various stacks from network-device manufacturers (i.e., switches and routers). In third place, there are 7M hosts with no label, which means Faults finds a matching signature for each of them, but Plata does not know what these devices are. The bottom half of the table, with a substantial count of cyber-physical systems and office equipment, is more alarming. These oftentimes run on default manufacturer passwords and allow reconfiguration using a built-in webserver. Investigating further, Table XV

TABLE XIV
TYPES OF DEVICES RUNNING WEBSERVERS

Device Type	Count	Fraction
General purpose	42,277,294	67%
Switch/router/gateway/network controller	8,854,290	14%
No label in database	7,038,785	11%
Printers	2,813,292	4.5%
RAID controller/NAS	1,348,895	2.1%
Video conferencing/telepresence	603,035	1.0%
Cyberphysical systems	91,033	0.14%
IP phones	61,400	0.10%

TABLE XV
UNPROTECTED INDUSTRIAL AND ENTERPRISE DEVICES

Device	Count	Type
Polycom HDX 8000 HD	266,565	Telepresence
Hickman ITV 450D	67,091	Telepresence
Cisco Unified IP Phone 7900 Series	27,151	IP Phone
AVTech RoomAlert/Rockwell Automation	21,756	Cyberphysical
Loytec L-DALI Lighting Control Systems	20,517	Cyberphysical
Codian Telepresence MCU	20,036	Telepresence
Polycom RealPresence Server 4000	18,977	Telepresence
AdTran IP Phone Manager	11,909	IP Phone
HWg-STE: Ethernet thermometer	11,826	Cyberphysical
D-Link DCS Series Internet Camera	9,279	Telepresence

TABLE XVI
OSes WITH EXPIRED SUPPORT LIFE CYCLES

OS	Count	Released
Windows 2000 / XP / 2003	1,512,725	2000 / 2001 / 2003
FreeBSD 7.3 / 8.0	433,978	2010 / 2009
Windows Server 2003 SP1 SP2	195,169	2005 / 2007
Windows Server 2000 SP4/XP SP3	146,421	2003 / 2008
FreeBSD 6.4	71,190	2008
Solaris 9 / Solaris 10	78,269	2003 / 2005
Mac OS X 10.4	36,834	2005
Windows 2000/XP SP1	9,623	2001 / 2002
Novell Netware OES 2 SP1	1,108	2005

shows the top-ten signatures from these categories, which include camera systems, building lighting controllers, and temperature monitors. They present high security risks to organizations because malicious actors may be able to use these systems to gain access to workplace audio/video recordings, printed documents, and environmental settings of critical infrastructure (e.g., cooling in data-centers).

With the recent leaks of NSA exploits and massive world-wide infection by ransomware WannaCry [24], [35], outdated operating systems (i.e., Windows XP/Server 2003) gained renewed attention. In Table XVI, we show several signatures that have reached the end of support and are no longer being patched to keep up with the latest vulnerabilities. These are obvious security threats; however, we find over 1.8M old Windows hosts still visible over the public Internet, 500K FreeBSD, and 78K Solaris. Faults not only allows for a timely measurement of such devices, but also paves the way for scalable, low-overhead Internet characterization, robust device identification, and better modeling of distortion θ experienced by the numerous hardware artifacts found on the Internet.

VIII. CONCLUSION

In this work, we developed novel theory and algorithms for improving OS-classification accuracy in single-probe

fingerprinting, measuring one-way Internet path properties, and extracting latent distributions of feature distortion. Simulations showed exceptional robustness of our EM techniques against various types of noise, as well as injection of unknown devices. Applied to Internet scans, this methodology can be used to detect vulnerable devices, as well as estimate stack popularity, network delays, packet loss, and header-tuning probabilities.

Future work involves construction of fingerprint databases with specimens that are pairwise separable under more complex distortion than just delay, detection of unknown stacks among the observations, automatic generation of signatures for them, and extensive comparison against nmap.

REFERENCES

- [1] H. J. Abdelnur, R. State, and O. Festor, "Advanced network fingerprinting," in *Proc. RAID*, Sep. 2008, pp. 372–389.
- [2] M. Antonakakis *et al.*, "Understanding the Mirai botnet," in *Proc. USENIX Secur.*, Aug. 2017, pp. 1093–1110.
- [3] Apple Support. *OS X Yosemite: Prevent Others From Discovering Your Mac*. Accessed: Oct. 2017. [Online]. Available: https://support.apple.com/kb/PH18642?locale=en_US
- [4] O. Arkin, "A remote active OS fingerprinting tool using ICMP," *USENIX login*, vol. 27, no. 2, pp. 14–19, Apr. 2002.
- [5] P. Auffret, "SinFP, unification of active and passive operating system fingerprinting," *J. Comput. Virol.*, vol. 6, no. 3, pp. 197–205, Aug. 2010.
- [6] T. Beardsley. (Jun. 2003). *Snacktime: A Perl Solution for Remote OS Fingerprinting*. [Online]. Available: <http://www.packetfu.com/wp/snacktime.html>
- [7] R. Beverly, "A robust classifier for passive TCP/IP fingerprinting," in *Proc. PAM*, Apr. 2004, pp. 158–167.
- [8] R. Beverly and A. Berger, "Server siblings: Identifying shared IPv4/IPv6 infrastructure via active fingerprinting," in *Proc. PAM*, Mar. 2015, pp. 149–161.
- [9] S. Boddy, R. Pompon, and R. Cohen. (Aug. 2019). The hunt for IoT: So easy to compromise, children are doing it. F5 Labs. [Online]. Available: <https://www.f5.com/labs/articles/threat-intelligence/the-hunt-for-iot-so-easy-to-compromise-children-are-doing-it>
- [10] Y.-C. Chen, Y. Liao, M. Baldi, S.-J. Lee, and L. Qiu, "OS fingerprinting and tethering detection in mobile networks," in *Proc. ACM IMC*, Nov. 2014, pp. 173–180.
- [11] H. K. J. Chu. (Jul. 2009). *Tuning TCP Parameters for the 21st Century*. [Online]. Available: <http://www.ietf.org/proceedings/75/slides/tcpm-1.pdf>
- [12] A. Crenshaw. (2008). *OSfuscate*. [Online]. Available: <http://www.irongeek.com/i.php?page=security/code>
- [13] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Statist. Soc.*, vol. 39, no. 1, pp. 1–38, 1977.
- [14] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman, "A search engine backed by Internet-wide scanning," in *Proc. ACM CCS*, Oct. 2015, pp. 542–553.
- [15] Z. Durumeric, E. Wustrow, and J. Halderman, "ZMap: Fast Internet-wide scanning and its security applications," in *Proc. USENIX Secur.*, Aug. 2013, pp. 605–620.
- [16] R. Ensafi, D. Fifield, P. Winter, N. Feamster, N. Weaver, and V. Paxson, "Examining how the great firewall discovers hidden circumvention servers," in *Proc. ACM IMC*, Oct. 2015, pp. 445–458.
- [17] X. Feng, Q. Li, H. Wang, and L. Sun, "Characterizing industrial control system devices on the Internet," in *Proc. IEEE ICNP*, Nov. 2016, pp. 1–10.
- [18] G. Shu and D. Lee, "Network protocol system fingerprinting—A formal approach," in *Proc. IEEE INFOCOM*, Apr. 2006, pp. 1–12.
- [19] H. O. Hartley, "Maximum likelihood estimation from incomplete data," *Biometrics*, vol. 14, no. 2, pp. 174–194, 1958.
- [20] J. Heidemann, Y. Pradkin, R. Govindan, C. Papadopoulos, G. Bartlett, and J. Bannister, "Census and survey of the visible Internet," in *Proc. ACM IMC*, Oct. 2008, pp. 169–182.
- [21] (Jul. 2020). *IRL OS-Fingerprinting Project*. [Online]. Available: <http://irl.cs.tamu.edu/projects/osf/>

- [22] Kaspersky Labs. (Jul. 2020). *Targeted Cyberattacks Logbook*. [Online]. Available: <https://apt.securelist.com>
- [23] M. Kearns, Y. Mansour, and A. Ng, "An information-theoretic analysis of hard and soft assignment methods for clustering," in *Proc. Uncertainty Artif. Intell.*, Aug. 1997, pp. 282–293.
- [24] Z. Kleinman. (May 2017). Cyber-attack: Is my computer at risk? BBC News. [Online]. Available: <http://www.bbc.com/news/technology-39896393>
- [25] T. Kohno, A. Broido, and K. C. Claffy, "Remote physical device fingerprinting," *IEEE Trans. Depend. Sec. Comput.*, vol. 2, no. 2, pp. 93–108, Feb. 2005.
- [26] E. Kollmann. *Chatter on the Wire: A Look at DHCP Traffic*. Accessed: Oct. 2017. [Online]. Available: <http://myweb.cableone.net/xnih/download/chatter-dhcp.pdf>
- [27] D. Leonard and D. Loguinov, "Demystifying service discovery: Implementing an Internet-wide scanner," in *Proc. ACM IMC*, Nov. 2010, pp. 109–122.
- [28] Z. Li, A. Goyal, Y. Chen, and V. Paxson, "Automating analysis of large-scale botnet probing events," in *Proc. ACM AsiaCCS*, Mar. 2009, pp. 11–22.
- [29] M. Luckie, R. Beverly, T. Wu, M. Allman, and K. Claffy, "Resilience of deployed TCP to blind attacks," in *Proc. ACM IMC*, Oct. 2015, pp. 13–26.
- [30] J. Matherly. (Jul. 2020). *Shodan Search Engine*. [Online]. Available: <https://shodan.io>
- [31] T. Matsunaka, A. Yamada, and A. Kubota, "Passive OS fingerprinting by DNS traffic analysis," in *Proc. IEEE AINA*, Mar. 2013, pp. 243–250.
- [32] C. McNab, *Network Security Assessment: Know Your Network*. Sebastopol, CA, USA: O'Reilly Media, 2007.
- [33] 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.
- [34] M. Technet. (Jul. 2012). *Stealth Mode in Windows Firewall with Advanced Security*. [Online]. Available: [https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/dd448557\(v=ws.10\)](https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/dd448557(v=ws.10))
- [35] M. Technet. (Oct. 2017). *Microsoft Security Bulletin MS17-010—Critical*. [Online]. Available: <https://docs.microsoft.com/en-us/security-updates/securitybulletins/2017/ms17-010>
- [36] A. Mirian *et al.*, "An Internet-wide view of ICS devices," in *Proc. IEEE PST*, Dec. 2016, pp. 96–103.
- [37] NetApplications. *Market Share Statistics for Internet Technologies*. Accessed: Oct. 2017. [Online]. Available: <http://netmarketshare.com/>
- [38] *Netcraft Web Server Survey*. Accessed: Oct. 2017. [Online]. Available: <http://news.netcraft.com/>
- [39] *Nmap*. Accessed: Oct. 2017. [Online]. Available: <http://nmap.org/>
- [40] G. Prigent, F. Vichot, and F. Harrouet, "IpMorph: Fingerprinting spoofing unification," *J. Comput. Virol.*, vol. 6, no. 4, pp. 329–342, Nov. 2010.
- [41] A. Quach, Z. Wang, and Z. Qian, "Investigation of the 2016 Linux TCP stack vulnerability at scale," in *Proc. ACM SIGMETRICS*, Jun. 2017, pp. 3:1–3:19.
- [42] G. Roulland and J.-M. Saffroy. *IP Personality*. Accessed: Oct. 2017. [Online]. Available: <http://ippersonality.sourceforge.net/>
- [43] A. Sarabi and M. Liu, "Characterizing the Internet host population using deep learning: A universal and lightweight numerical embedding," in *Proc. ACM IMC*, Oct. 2018, pp. 133–146.
- [44] S. Shah. (May 2004). *An Introduction to HTTP Fingerprinting*. http://net-square.com/httpprint_paper.html
- [45] Z. Shamsi and D. Loguinov, "Unsupervised clustering under temporal feature volatility in network stack fingerprinting," in *Proc. ACM SIGMETRICS*, Jun. 2016, pp. 127–138.
- [46] Z. Shamsi, A. Nandwani, D. Leonard, and D. Loguinov, "Hershel: Single-packet OS fingerprinting," in *Proc. ACM SIGMETRICS*, Jun. 2014, pp. 195–206.
- [47] Z. Shamsi, D. B. H. Cline, and D. Loguinov, "Faulds: A non-parametric iterative classifier for Internet-wide OS fingerprinting," in *Proc. ACM CCS*, Oct. 2017, pp. 971–982.
- [48] Z. Shamsi, D. B. Cline, and D. Loguinov, "Faulds: A non-parametric iterative classifier for Internet-wide OS fingerprinting," Texas A&M Univ., College Station, TX, USA, Tech. Rep. 2017-8-2, Aug. 2017. [Online]. Available: <http://irl.cs.tamu.edu/publications/>
- [49] U. Shankar and V. Paxson, "Active mapping: Resisting NIDS evasion without altering traffic," in *Proc. IEEE SP*, May 2003, pp. 44–61.
- [50] B. Skaggs, B. Blackburn, G. Manes, and S. Sheno, "Network vulnerability analysis," in *Proc. IEEE MWSCAS*, Aug. 2002, pp. 493–495.
- [51] M. Smart, G. R. Malan, and F. Jahanian, "Defeating TCP/IP stack fingerprinting," in *Proc. USENIX Secur.*, Jun. 2000, pp. 229–240.
- [52] A. K. Sood and R. J. Enbody, "Targeted Cyberattacks: A superset of advanced persistent threats," *IEEE Security Privacy*, vol. 11, no. 1, pp. 54–61, Jan. 2013.
- [53] G. Taleck, "Ambiguity resolution via passive OS fingerprinting," in *Proc. RAID*, Sep. 2003, pp. 192–206.
- [54] G. Taleck, "SYNSCAN: Towards complete TCP/IP fingerprinting," in *Proc. CanSecWest*, Apr. 2004, pp. 1–12.
- [55] F. Veysset, O. Courtay, and O. Heen. (Apr. 2002). *New Tool and Technique for Remote Operating System Fingerprinting*. [Online]. Available: <http://www.ouah.org/ring-full-paper.pdf>
- [56] K. Wang. (2004). *Frustrating OS Fingerprinting With Morph*. [Online]. Available: <http://hackerpoetry.com/images/defcon-12/dc-12-presentations/Wang/dc-12-wang.pdf>
- [57] K. Yang, Q. Li, and L. Sun, "Towards automatic fingerprinting of IoT devices in the cyberspace," *Comput. Netw.*, vol. 148, pp. 318–327, Jan. 2019.
- [58] F. V. Yarochkin, O. Arkin, M. Kydyraliev, S.-Y. Dai, Y. Huang, and S.-Y. Kuo, "Xprobe2++: Low volume remote network information gathering tool," in *Proc. IEEE/IFIP DSN*, Jun. 2009, pp. 205–210.
- [59] M. Zalewski. (Apr. 2001). *Strange Attractors and TCP/IP Sequence Number Analysis*. [Online]. Available: <http://lcamtuf.coredump.cx/newtcp/>
- [60] M. Zalewski. (2012). *p0f v3: Passive Fingerprinter*. [Online]. Available: <http://lcamtuf.coredump.cx/p0f3>



Zain Shamsi received the B.S. degree (Hons.) in computer science from Trinity University, San Antonio, TX, USA, in 2008, and the Ph.D. degree in computer science from Texas A&M University, College Station, TX, USA, in 2017. He currently works with Applied Research Laboratories (ARL), Austin, TX, USA. His research interests include security, network protocols, and Internet measurement.



Daren B. H. Cline received the B.S. degree in mathematics from the Harvey Mudd College, Claremont, CA, USA, in 1978, and the M.S. and Ph.D. degrees in statistics from Colorado State University, Fort Collins, CO, USA, in 1980 and 1983, respectively. In 1984, he joined the Department of Statistics, Texas A&M University, College Station, TX, USA, where he is currently a Tenured Professor. His current research interests include applied stochastic processes, Markov chain theory, heavy-tailed distributions, and nonlinear time series models.



Dmitri Loguinov (Senior Member, IEEE) received the B.S. degree (Hons.) in computer science from Moscow State University, Russia, in 1995, and the Ph.D. degree in computer science from The City University of New York in 2002. He is currently a Professor with the Department of Computer Science and Engineering, Texas A&M University, College Station. His research interests include P2P networks, information retrieval, congestion control, and Internet measurement and modeling.