

The Mason Test: A Defense Against Sybil Attacks in Wireless Networks Without Trusted Authorities

Yue Liu, David R. Bild, *Member, IEEE*, Robert P. Dick, *Member, IEEE*, Z. Morley Mao, and Dan S. Wallach

Abstract—Wireless networks are vulnerable to Sybil attacks, in which a malicious node poses as many identities in order to gain disproportionate influence. Many defenses based on spatial variability of wireless channels exist, but depend either on detailed, multi-tap channel estimation—something not exposed on commodity 802.11 devices—or valid RSSI observations from multiple trusted sources, e.g., corporate access points—something not directly available in ad hoc and delay-tolerant networks with potentially malicious neighbors. We extend these techniques to be practical for wireless ad hoc networks of commodity 802.11 devices. Specifically, we propose two efficient methods for separating the valid RSSI observations of behaving nodes from those falsified by malicious participants. Further, we note that prior signalprint methods are easily defeated by mobile attackers and develop an appropriate challenge-response defense. Finally, we present the Mason test, the first implementation of these techniques for ad hoc and delay-tolerant networks of commodity 802.11 devices. We illustrate its performance in several real-world scenarios.

Index Terms—Wireless networks, security, ad hoc networks, Sybil attack.



1 INTRODUCTION

THE open nature of wireless ad hoc networks (including delay-tolerant networks [1]) enables applications ranging from collaborative environmental sensing [2] to emergency communication [3], but introduces numerous security concerns since participants are not vetted. Solutions generally rely on a majority of the participants following a particular protocol, an assumption that often holds because physical nodes are expensive. However, this assumption is easily broken by a Sybil attack. A single physical entity can pretend to be multiple participants, gaining unfair influence at low cost [4]. Newsome et al. survey Sybil attacks against various protocols [5], illustrating the need for a practical defense.

Proposed defenses (see Levine et al. for a survey [6]) fall into two categories. *Trusted certification* methods [7], [8] use a central authority to vet potential participants and thus are not useful in open ad hoc (and delay-tolerant) networks. *Resource testing* methods [9], [10], [11], [12] verify the resources (e.g., computing capability, storage capacity, real-world social relationships, etc.) of each physical entity. Most are easily defeated in ad hoc networks of resource-limited mobile devices by attackers with access to greater resources, e.g., workstations or data centers.

One useful class of defenses is based on the natural

spatial variation in the wireless propagation channel, an implicit resource. Channel responses are uncorrelated over distances greater than half the transmission wavelength [13] (6 cm for 2.4GHz 802.11), so two transmissions with the same channel response are very likely to be from the same location and device [14], [15]. Note that two transmitters may be close enough, i.e., ~ 6 cm, to produce the same channel response, but this case is rare in practice¹. One class of Sybil defenses based on this observation uses specialized hardware to accurately measure and compare channel responses [15]. However commodity devices are not equipped with such hardware.


Commodity devices expose an aggregate, scalar value, the received signal strength. RSSI can be changed by varying transmit power, so a vector of observations from multiple receivers—a *signalprint*—is used instead. Several authors have proposed such methods [16], [17], [18], [19], [20], [21] assuming trusted, true observations from, for example, access points (Figure 1a). In open ad hoc networks, observations are untrusted, coming from potentially lying neighbors (Figure 1b). In this case observations falsified by attackers can lead to incorrect conclusions (Figure 1c). Trust-less methods have been proposed, but have various limitations (e.g., devices must have uniform transmit power [22] or the method may be used only in an outdoor environments with predictable propagation ranges [23]). Instead, a general method to separate true and false observations is needed.

We observe that with high probability attackers cannot produce false observations that make conforming

- Y. Liu, D. R. Bild, R. P. Dick, and Z. M. Mao are with the Department of Electrical and Computer Engineering, University of Michigan, Ann Arbor, MI 48109.
E-mail: {liuyue, drbild, dickrp, zmao}@umich.edu
- D. S. Wallach is with the Department of Computer Science, Rice University, Houston, TX 77005.
E-mail: dwallach@cs.rice.edu

This work was supported, in part, by NSF under award TC-0964545.

1. In our experiments with smartphone users, distinct transmitters displayed similar channel responses in fewer than 0.01% of cases (see Figure 15 in Section 9).



(a) RSSI observations from trusted APs identify the Sybils, S_1 and S_2 , presented by attacker M .

(b) In ad hoc networks, the participants themselves act as observers, but can maliciously report falsified values.

(c) If I believes the falsified observations from S_1 and S_2 , it will incorrectly accept them and reject A and B as Sybil.

Fig. 1. Prior work [15], [16] assumes trusted RSSI observations, not generally available in ad hoc and delay-tolerant networks. We present a technique for a participant to separate true and false observations, enabling use in ad hoc networks. (Arrows point from transmitter to observer.)

identities look Sybil, due to the unpredictability of the wireless channels. We exploit this weakness to bound the number of misclassified identities. In cases when conforming nodes outnumber physical attacking nodes (a major motivating factor for the Sybil attack), we develop a notion of consistency that enables fully-accurate classification.

Most past work assumes nodes are stationary, as moving attackers can easily defeat naïve signalprint-based detection. Although reasonable for conforming nodes, e.g., most human-carried smartphones are stationary over short time-spans, this is too strong an assumption for attackers. We remove this restriction on the attack model and defeat moving attacks by detecting mobile nodes. Moving nodes are not permitted to participate in other protocols because their “Sybilness” cannot be verified. They can be tested again when stationary. Xiao et al. noted that successive transmissions from the same stationary node should have the same signalprint and attackers likely cannot quickly (i.e., in milliseconds) switch between precise positions [15], but did not further develop or evaluate a method making use of this observation. We develop a challenge–response protocol from this idea and study its performance on real deployments.

On a high level, we seek to allow a wireless network participant to occasionally determine which of its one-hop neighbors are non-Sybil. Verified non-Sybil participants, uniquely identified by their public keys, may safely participate in other protocols. In mobile networks, the process must be repeated occasionally (e.g., once per hour) as the topology changes. Safety is more important than system performance, so nearly all Sybil identities must be detected. In most applications it is acceptable for some non-Sybils to be rejected, e.g., any that were moving during the test.

We make the following primary contributions.

- We design two methods of $O(n^3)$ complexity to separate true and false RSSI observations, enabling signalprint-based Sybil detection in ad hoc networks of nodes without mutual trust. The first method gives partial separation, bounding the number of misclassified identities. The second provides full separation, but works only when conforming nodes outnumber physical attacking nodes.
- We prove conditions under which a participant can fully separate true and false observations.
- We develop a challenge-response protocol to detect attackers attempting to use motion to defeat the signalprint-based Sybil defense.
- We describe the Mason test, a practical protocol for Sybil defense based on these ideas. We implemented the Mason test as a Linux kernel module for 802.11 ad hoc networks² and characterize its performance in real-world scenarios.

2 RELATED WORK

Many Sybil defense techniques are built on wireless channel-based resource testing, because placing transmitters in many locations is much more difficult than acquiring additional computation or memory resources. Xiao et al. observe that in OFDM-based 802.11, the coherence bandwidth is much smaller than the system bandwidth and thus the channel response estimates at well-spaced frequency taps are uncorrelated, forming a vector unique to the transmitter location and robust to changes in transmitter power [15]. Li et al. use both temporal and multiple-tone channel probing to profile

2. <http://github.com/EmbeddedAtUM/mason/>

transmissions and detect anomalous behaviors [19]. Unfortunately commodity 802.11 devices do not expose these detailed channel estimates to the driver and operating system, restricting this technique to specialized hardware or firmware.

Faria et al. and Demirbas et al. independently developed the RSSI-based signalprint technique to greatly simplify channel estimations while maintaining high classification performance [16], [17]. This class of work [16], [17], [18], [19], [20], [21] has two disadvantages. First they rely on trusted external measurements, e.g., observations by trusted 802.11 access points, which are generally unavailable in open ad hoc networks. Our work builds on their ideas, but removes reliance on any external measurements. Second, their attack models are restricted to stationary devices, despite the ease of deploying moving attackers. Our work detects and rejects moving nodes, instead of accepting them as non-Sybil.

Lv et al. developed a method based on one-dimensional signalprints, which therefore does not rely on external measurements [22]. However, it assumes a uniform transmit power for all devices, including attacking devices.

Bouassida et al. developed a trust-less method for vehicular area networks. Instead of relying on external measurements, the “verifier” obtains uncorrelated measurements by changing its own reception locations and uses these measurements to locate the transmitter. It also rejects moving nodes with significant location changes over multiple measurements [23]. However, this technique relies on a predictable propagation model that fails to capture the notorious variations of wireless channels. Our method does not assume any propagation model. Instead, we rely on the unpredictability of wireless signal propagation to defeat lying attackers.

3 PROBLEM FORMULATION AND BACKGROUND

In this section, we define our problem, summarize the solution framework, describe our attack model, and briefly review the signalprint method.

3.1 Problem Formulation

Our goal is to extend signalprint-based Sybil detection methods to work without a priori trust in any observers, allowing any participant in an open wireless network to determine which of its one-hop neighbors are non-Sybil. The solution framework is illustrated in Figure 2 with five participants. We assume an arbitrary identity (or condition) starts the process. Participants first take turns broadcasting a probe packet while all others record the observed RSSIs (Figure 2a). These observations are then shared, although malicious nodes may lie. Figure 2b shows every participant after this exchange, with observations from all five participants. Finally each participant individually selects a (hopefully truthful) subset of observers for signalprint-based Sybil classification (Figure 2c).

This paper presents our method for truthful subset selection and fleshes out this framework into a usable, safe, and secure protocol. As with any system intended for real-world use, we had to carefully balance system complexity and potential security weaknesses. Section 10 discusses these choices and related potential concerns.

3.2 Attack Model

We model attackers who operate commodity devices, but not specialized hardware. Commodity devices can be obtained in large scale by compromising those owned by normal network participants, a more practical attack vector than distributing specialized hardware at the same scale. Specifically, we assume attackers have the following capabilities and restrictions.

- 1) Attackers may collude through arbitrary side channels.
- 2) Attackers may accumulate information, e.g., RSSIs, across multiple rounds of the Mason test.
- 3) Attackers have limited ability to predict the RSSI observations of other nodes, e.g., 7 dBm uncertainty (see Section 6), precluding fine-grained pre-characterization.
- 4) Attackers can control transmit power for each packet, but not precisely or quickly steer the output in a desired direction, i.e., they are not equipped for antenna array-based beam-forming.
- 5) Attackers can move their devices, but cannot quickly and precisely switch them between multiple positions, e.g., they do not have high-speed, automated electromechanical control.

One common denial-of-service (DOS) attack in wireless networks—jamming the channel—cannot be defended against by commodity devices. Thus, we do not defend against other more-complicated DOS attacks. However, note that ad hoc and delay-tolerant networks are much more resistant than infrastructured networks to such attacks, because a single attack can affect only a small portion of the network.

Notably, we assume attackers do not have per-antenna control of MIMO (Multiple-Input and Multiple-Output) [24] devices. Such control would defeat the signalprint method (even with trusted observers), but is not a feasible attack. Commodity MIMO devices (e.g., 802.11n adapters) do not expose this control to software and thus are not suitable attack vectors. Distributing specialized MIMO-capable hardware over large portions of the network would be prohibitively expensive.

We believe that the signalprint method can be extended to MIMO systems (see our technical report for an overview [25]), but doing so is beyond the scope of this work. Our focus is extending signalprint-based methods to ad hoc networks of commodity devices by removing the requirement for trusted observations.

3.3 Review of Signalprints

We briefly review the signalprint method. See prior work for details [15], [17]. A *signalprint* is a vector of RSSIs

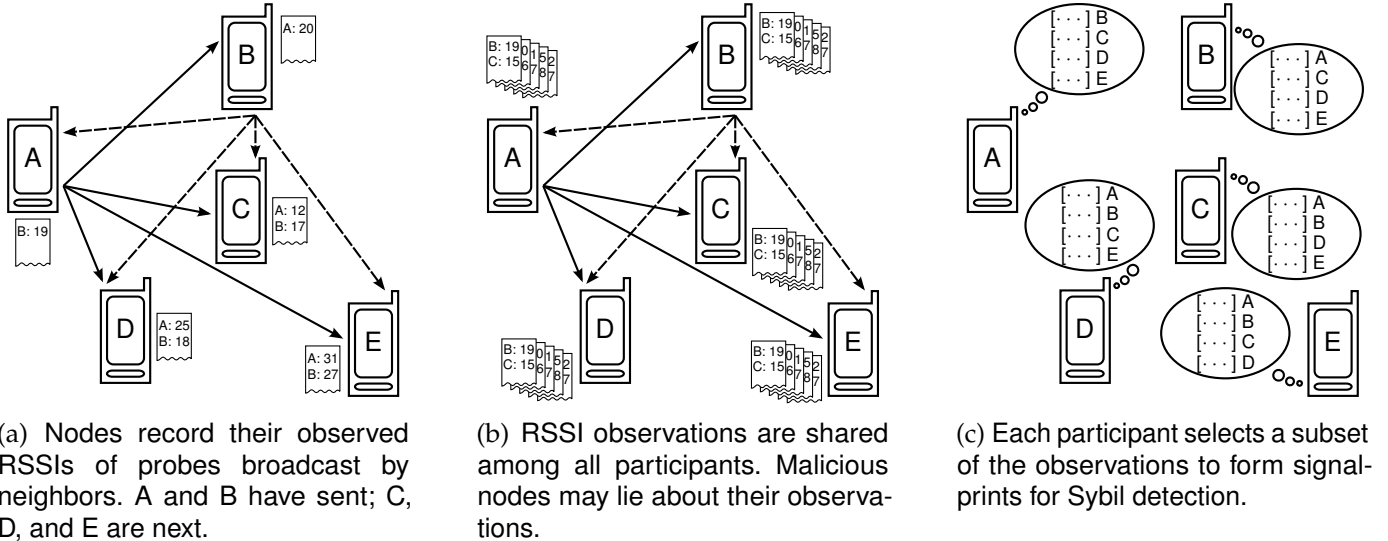


Fig. 2. The solution framework for signalprint-based Sybil detection in ad hoc networks. This paper fleshes out this concept into a safe and secure protocol, the Mason test.

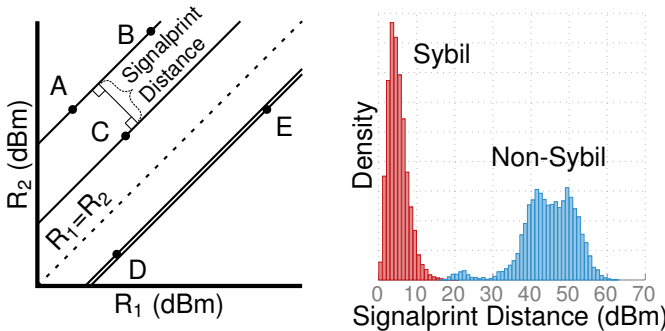


Fig. 3. Sybils, A – B and D – E , occupy nearby slope-1 lines.

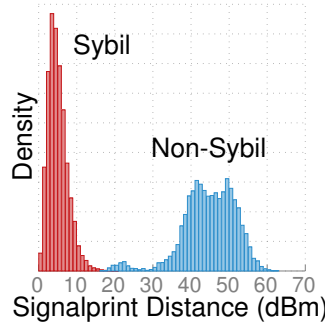


Fig. 4. The distance threshold trades false positives for negatives.

at multiple observers for a single transmission. Ignoring noise, the vector of received powers (in logarithmic units, e.g., dBm) at multiple receivers for a given transmission can be modeled [13] as $\vec{s} = \vec{h} + p\vec{1}$, where p is the transmit power and \vec{h} is the attenuation vector, a function of the channel amplitude response and the receiver characteristics. Transmissions from different locations have uncorrelated signalprints, as the channel responses are likely uncorrelated. Those from the same location, however, share a channel response and will be correlated. That is, for two transmissions a and b from the same location with transmit powers p_a and $p_b = p_a + c$, the signalprints $\vec{s}_b = \vec{h} + p_a\vec{1}$ and $\vec{s}_c = \vec{h} + (p_a + c)\vec{1}$ are related as $\vec{s}_c = \vec{s}_a + c\vec{1}$. In other words, all observers see the same RSSI difference c for the two transmissions.

This is illustrated geometrically in Figure 3 for a two-receiver signalprint. A and B are Sybil, while C is not. D and E are also Sybil, but due to noise the signalprints are not perfectly correlated. Instead, signalprints on lines closer than some threshold are taken to be Sybil.

Definition. The *signalprint distance* $d(\vec{s}_a, \vec{s}_b)$ between two

signalprints \vec{s}_a and \vec{s}_b is the perpendicular distance between the slope-1 lines containing them. Letting

$$\vec{w} \triangleq \vec{s}_a - \vec{s}_b$$

be the distance vector between the signalprints and

$$\vec{v}_\perp \triangleq \vec{w} - \frac{\vec{w} \cdot \vec{1}}{\|\vec{1}\|^2} \vec{1}$$

be the vector rejection of \vec{w} from $\vec{1}$, then

$$d(\vec{s}_a, \vec{s}_b) = \|\vec{v}_\perp\|.$$

Figure 4 shows the distance distributions for Sybil and non-Sybil identities using measurement data for commodity Android devices (size-4 signalprints from the “Outdoor” experiment in Section 9). The two distributions are well separated with small overlap, so the threshold choice trades false positives for false negatives. A good threshold supports detection of at least 99.9% of Sybils while accepting at least 95% of non-Sybils, as reported by previous research [15], [17] and confirmed by our own measurement (see Figure 13).

4 SYBIL CLASSIFICATION FROM UNTRUSTED SIGNALPRINTS

In this section we describe two methods to detect Sybil identities using untrusted RSSI observations. In both cases, a set of candidate views containing the true view (with high probability) is generated. The accepted view is chosen by a view selection policy. The first method selects the view indicting the most Sybils, limiting the total number of incorrect classifications. The second selects the true view, but works only when conforming nodes outnumber physical attacker nodes.

4.1 The Limited Power of Falsified Observations

Our key observation is that falsified RSSI observations have limited power. Although falsifying observations to make Sybil identities look non-Sybil is easy, it is extremely difficult to make a non-Sybil look Sybil. To see this, recall that a pair of identities is considered Sybil only if all observers report the same RSSI difference for the pair's transmissions. For any pair of identities, the initiator observes an RSSI difference itself. On one hand, making true Sybils appear non-Sybil is easy, because randomly chosen values almost certainly fail to match this difference. Making a non-Sybil look Sybil, however, requires learning the difference observed by the initiator itself, which is kept secret. Guessing is difficult due to the unpredictability of the wireless channels. Our methods rely on this difficulty. They are developed formally in the rest of this section. Quantitative characterizations are described in Section 6. To summarize, the success probability for a guessing attacker is less than one percent in common situations, i.e., when conforming nodes outnumber physical attackers by more than $1.5\times$.

4.2 Terminology

I is the set of participating identities. Each is either Sybil or non-Sybil and reports either true or false³ RSSI observations, partitioning the identities by their Sybilness (sets S and NS) and the veracity of their reported observations (sets T and L).

	S	NS
L	LS	LNS
T	TS	C

Truth-telling, non-Sybil identities are called *conforming* (set C). Liars and Sybil identities are called *attacking* (sets LS , LNS , and TS). Our goal is to distinguish the S and NS partitions using the reported RSSI observations without first knowing the L and T partitions.

Definition. An *initiator* is the node performing Sybil classification.⁴ It trusts its own RSSI observations, but not others.

Definition. A *receiver set*, denoted by R , is a subset of identities ($R \subseteq I$) whose reported RSSI observations, combined with the initiator's, form signalprints. Those with liars ($R \cap L \neq \emptyset$) produce incorrect classifications and those with only truth-tellers ($R \subseteq T$) produce the correct classification.

Definition. A *view*, denoted by V , is a classification of identities as Sybil and non-Sybil. Those classified as Sybil

(non-Sybil) are said to be Sybil (non-Sybil) *under* V and are denoted by the subset V_S (V_{NS}). A view V obtained from the signalprints of a receiver set R is *generated* by R , denoted by $R \mapsto V$ (read: R generates V), and can be written $V(R)$. Identities in R are considered non-Sybil, i.e., $R \subseteq V_{NS}(R)$. A *true view*, denoted by \bar{V} , correctly labels all identities, i.e., $\bar{V}_S = S$ and $\bar{V}_{NS} = NS$. Similarly, a *false view*, denoted by \hat{V} , incorrectly labels some identities, i.e., $\hat{V}_S \neq S$ and $\hat{V}_{NS} \neq NS$.

Definition. Incorrectly labeling non-Sybil identities as Sybil is called *collapsing*.

Assumption. To clearly illustrate the impact of intentionally-falsified observations, we first assume that true RSSI observations are noise-free and thus always generate the true view. In Section 4.7, we extend the method to handle real-world observations containing, for example, random noise and discretization error.

4.3 Approach Overview

A general separation method does not exist, because different scenarios can lead to the same reported RSSI observations. To illustrate, consider identities $I = \{A|B\}$ reporting observations such that

$$R \subseteq A \mapsto V^1 = \{V_{NS}^1 = A | V_S^1 = B\} \text{ and} \\ R \subseteq B \mapsto V^2 = \{V_{NS}^2 = B | V_S^2 = A\}$$

and two different scenarios x and y such that

$$\text{in } x, \{T^x = A | L^x = B\} = I \text{ and} \\ \text{in } y, \{T^y = B | L^y = A\} = I.$$

$R \subseteq T \mapsto \bar{V}$, so V^1 and V^2 are both true views, the former in scenario x and the latter in scenario y . In other words, identities in A could be Sybil (as claimed by B) or those in B could be Sybil (as claimed by A). Either view could be correct; it depends on which group is lying. Consequently, no method can always choose the correct view.

We instead develop two different approaches. The first method, the *maximum Sybil policy*, simply bounds the number of misclassified identities by selecting the view reporting the most Sybils. This selected view must indict at least as many as the true view, bounding the accepted Sybils by the number of collapsed conforming identities. Collapsing is difficult, limiting the incorrect classifications.

The second method, the *view consistency policy*, allows complete separation, but requires that the following conditions be met.

- All views correctly classify some conforming identities (likely true because collapsing identities is difficult).
- Conforming identities outnumber lying, non-Sybils (a major motivating factor for the Sybil attack).

3. A reported RSSI observation is considered *false* if some signalprints containing it misclassify some identities.

4. All participants perform classification individually, so each is the initiator in its own classification session.

TABLE 1
Definitions of Terms and Symbols

	Definition	Notes
Sets of Identities		
I	all participating identities	
NS	all non-Sybil identities	
S	all Sybil identities	$I = \{NS S\}$
T	all truthful identities	
L	all lying identities	$I = \{T L\}$
C	all conforming, or truthful, non-Sybil, identities	$NS = \{C LNS\}$
LNS	all lying, non-Sybil identities	$S = \{TS LS\}$
TS	all truthful, Sybil identities	$T = \{C TS\}$
LS	all lying, Sybil identities	$L = \{LNS LS\}$
V_{NS}	all identities labeled non-Sybil by view V	
V_S	all identities labeled Sybil by view V	$I = \{V_{NS} V_S\}$
R (receiver set)	identities used to form signalprints	
Views		
V (view)	a Sybil–non-Sybil labeling of I	
\bar{V} (true view)	a view that correctly labels all identities	$\bar{V}_{NS} = NS$ and $\bar{V}_S = S$
\hat{V} (false view)	a view that incorrectly labels some identities	$\hat{V}_{NS} \neq NS$ and $\hat{V}_S \neq S$
$V(R)$	the view generated by receiver set R	
Terms		
<i>generates</i> ($R \mapsto V$)	a receiver set generates a view	
<i>initiator</i>	node performing the Sybil classification	
<i>collapse</i>	classify a non-Sybil identity as Sybil	

This approach follows from the idea that true observations are trivially self-consistent, while lies often contradict themselves. We develop a notion of consistency that allows separation of true and false observations.

4.4 Maximum Sybil Policy: Select the View Claiming the Most Sybil Identities

In this section, we prove that the maximum Sybil policy—selecting the view claiming the most Sybil identities—produces a classification with bounded error. The number of incorrectly-accepted Sybil identities is bounded by the number of collapsed conforming identities.

Lemma 1. *The selected view V claims at least as many Sybil identities as actually exist, i.e., $|V_S| \geq |S|$.*

Proof: Since the true view \bar{V} claiming $|S|$ Sybils always exists, the selected view can claim no fewer. \square

Theorem 1. *The selected view V misclassifies no more Sybil identities than it collapses conforming identities, i.e., $|V_{NS} \cap S| \leq |V_S \cap NS|$.*

Proof: Intuitively, claiming the minimum $|S|$ Sybil identities requires that each misclassified Sybil be compensated for by a collapsed non-Sybil identity. Formally, combining $|V_S \cup V_{NS}| = |S \cup NS|$ with Lemma 1 yields $|(V_S \cup V_{NS}) \cap S| \leq |(S \cup NS) \cap V_S|$. Removing the common $V_S \cap S$ from both sides gives $|V_{NS} \cap S| \leq |V_S \cap NS|$. \square

Theorem 1 bounds the misclassifications by the attacker’s collapsing power, $|V_S \cap NS|$. Although $|V_S \cap NS|$ is small (see Section 6), one Sybil is still accepted for each conforming identity collapsed. The next few sections develop a second method that allows accurate classification, but only when conforming nodes outnumber attackers.

4.5 View Consistency Policy: Selecting \bar{V} if $LNS = \emptyset$

Our view consistency policy stems from the intuition that lies often contradict each other. It is introduced here using the following unrealistic assumption, which we remove in Section 4.6.

Restriction 1. All liars are Sybil, i.e., $LNS = \emptyset$, and thus all non-Sybil identities are truthful, i.e., $NS \subseteq T$.

Restriction 1 endows the true view with a useful property: all receiver sets comprising the non-Sybil identities under the true view will generate the true view. We formalize this notion of consistency as follows.

Definition. A view is *view-consistent* if and only if all receiver sets comprising a subset of the non-Sybil identities under that view generate the same view, i.e., V is view-consistent iff $\forall R \in 2^{V_{NS}} : R \mapsto V$.

Lemma 2. *Under Restriction 1, the true view is view-consistent, i.e., $\forall R \in 2^{\bar{V}_{NS}} : R \mapsto \bar{V}$.*

Proof: Consider the true view \bar{V} . By definition, $\bar{V}_{NS} = NS$. By Restriction 1, $NS \subseteq T$ and thus, $\bar{V}_{NS} \subseteq T$. $\forall R \in 2^T \mapsto \bar{V}$, so $\forall R \in 2^{\bar{V}_{NS}} : R \mapsto \bar{V}$. \square

Were all false views not consistent, then consistency could be used to identify the true view. However, a fully omniscient attacker could theoretically generate a false, consistent view by collapsing all conforming identities. In practice, the difficulty of collapsing identities prevents this. We formalize this attacker limitation as follows.

Condition 1. All receiver sets correctly classify at least one conforming identity, i.e., $\forall R \in 2^I : V_{NS}(R) \cap C \neq \emptyset$.

Justification: Collapsing conforming identities re-

quires knowing the hard-to-predict initiator's RSSI observations. Section 6 quantifies the probability that this condition holds.

Lemma 3. *Under Condition 1, a view generated by a receiver set containing a liar is not view-consistent, i.e., $R \cap L \neq \emptyset$ implies $V(R)$ is not view-consistent.*

Proof: Consider such a receiver set R with $R \cap L \neq \emptyset$. By Condition 1, $r \triangleq V_{\text{NS}}(R) \cap C$ is not empty and since $r \subseteq C \subseteq T$, $r \mapsto \bar{V}$. By the definition of a liar, $V(R) \neq \bar{V}$ and thus R is not consistent. \square

Theorem 2. *Under Restriction 1 and Condition 1 and assuming $C \neq \emptyset$, exactly one consistent view is generated across all receiver sets and that view is the true view.*

Proof: By Lemma 2 and Lemma 3, only the true view is consistent, so we need only show that at least one receiver set generates the true view. $C \neq \emptyset$ and thus $R = C \mapsto \bar{V}$. \square

This result suggests a method to identify the true view—select the only consistent view. Restriction 1 does not hold in practice, so we develop methods to relax it.

4.6 Achieving Consistency by Eliminating LNS

Consider a scenario with some non-Sybil liars. The true view would be consistent were the non-Sybil liars excluded from consideration. Similarly, a false view could be consistent were the correctly classified conforming identities excluded. If the latter outnumber the former, this yields a useful property: the consistent view over the largest subset of identities, i.e., that with the fewest excluded, is the true view, as we now formalize and prove.

Condition 2. The number of conforming identities is strictly greater than the number of non-Sybil liars, i.e., $|C| > |LNS|$.

Justification: This is assumed by networks whose protocols require a majority of the nodes to conform. In others, it may hold for economic reasons—deploying as many nodes as the conforming participants is expensive.

Condition 3. Each receiver set either correctly classifies at least $|LNS|+1$ conforming identities as non-Sybil or the resulting view, when all correctly classified conforming identities are excluded, is not consistent, i.e., $\forall R \in 2^I : (|V_{\text{NS}}(R) \cap C| \geq |LNS| + 1) \vee (\exists Q \in 2^{V_{\text{NS}}(R) \setminus C} : V(Q) \neq V(R))$. Note that this implies Condition 2.

Justification: This is an extension of Condition 1. Section 6 quantifies the probability that it holds.

Lemma 4. *Under Condition 2 and Condition 3, the largest subset of I permitting a consistent view is $I \setminus LNS$.*

Proof: $I \setminus LNS$ permits a consistent view, per Lemma 2. Let $E_R \triangleq \hat{V}_{\text{NS}}(R) \cap C$ be the set of correctly classified conforming nodes for a lying receiver set R , i.e., $R \cap L \neq \emptyset$.

$I \setminus E_R$ is the largest subset possibly permitting a consistent view under R . By Condition 3, $\forall R : |E_R| \geq |LNS| + 1$. \square

Theorem 3. *Under Condition 2 and Condition 3, the largest subset of I permitting a consistent view permits just one consistent view, the true view.*

Proof: This follows directly from Lemma 4 and Theorem 2. \square

In the next section, we extend the approach to handle the noise inherent to real-world signalprints.

4.7 Extending Consistency to Handle Noise

Noise prevents true signalprints from always generating the true view. Observing from prior work that the misclassifications are bounded (e.g., more than 99% of Sybils detected with fewer than 5% of conforming identities collapsed [15], [17]), we extend the notion of consistency as follows.

Definition. Let γ_n be the maximum fraction⁵ of non-Sybil identities misclassified by a size- n receiver set. Prior work suggests $\gamma_4 = 0.05$ is appropriate (for $|C| > 20$) [15], [17].

Definition. A view is γ_n -consistent if and only if all size- n receiver sets that are subsets of the non-Sybil identities under that view generate a γ_n -similar view. Two views V^1 and V^2 are γ_n -similar if and only if

$$\left(\frac{|V_{\text{NS}}^1 \cap V_{\text{NS}}^2|}{|V_{\text{NS}}^1 \setminus V_{\text{NS}}^2|} > \frac{1 - 2\gamma_n}{\gamma_n} \right) \wedge \left(\frac{|V_{\text{NS}}^1 \cap V_{\text{NS}}^2|}{|V_{\text{NS}}^2 \setminus V_{\text{NS}}^1|} > \frac{1 - 2\gamma_n}{\gamma_n} \right)$$

This statement captures the intuitive notion that V_{NS}^1 and V_{NS}^2 should contain the same identities up to differences expected under the γ_n bound. A view is γ_n -true if it is γ_n -similar to the true view.

Lemma 5. *Under Restriction 1, the view generated by any truthful receiver set of size n is γ_n -consistent.⁶*

Proof: Consider two views V^1 and V^2 generated by conforming receiver sets. Each correctly classifies at least $(1 - \gamma_n)$ of the non-Sybil identities, so $|V_{\text{NS}}^1 \cap V_{\text{NS}}^2| \geq (1 - 2\gamma_n)|NS|$. Each misclassifies at most γ_n of the non-Sybil identities, so $|V_{\text{NS}}^1 \setminus V_{\text{NS}}^2| \leq \gamma_n|NS|$ and similar for $V_{\text{NS}}^2 \setminus V_{\text{NS}}^1$. The ratio of these bounds is the result. \square

Substituting γ -consistency for pure consistency, Condition 3 still holds with high (albeit different) probability, quantified in Section 6. An analogue of Theorem 3 follows.

Theorem 4. *Under Condition 3, the γ_n -consistent view of the largest subset of I permitting such a view is γ_n -true.*

In Section 5 we describe an efficient algorithm to identify the largest subset permitting a γ -consistent view and thus the correct (up to errors expected due to signalprint noise) Sybil classification.

⁵ γ_n is an upper bound on the total fraction misclassified, not the probability that an individual identity is misclassified.

⁶ This assumes that the false negative bound is negligible. If it is not, a similar notion of γ, σ -consistency, where σ is the false negative bound, can be used. In practice σ is quite small [15], [17], so simple γ_n -consistency is fine.

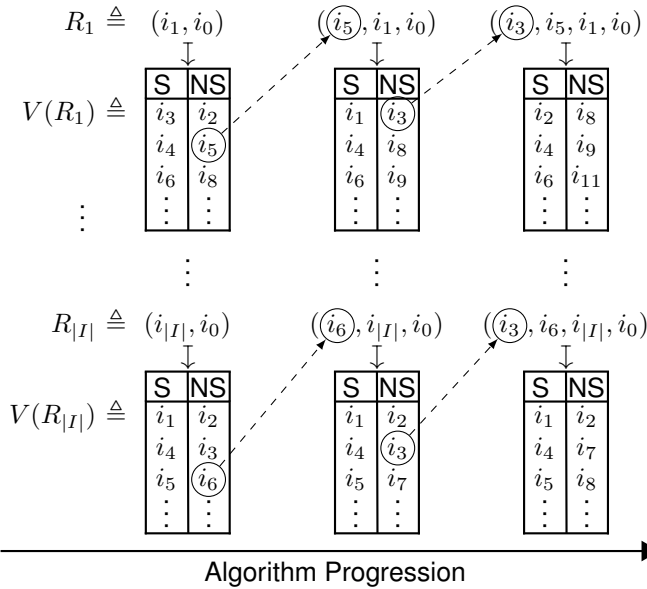


Fig. 5. Illustration of Algorithm 1. All $|I|$ size-2 receiver sets are increased to size-4 by iteratively adding a random identity from those labeled non-Sybil by the current set. With high probability, at least one of the final sets will contain only conforming identities.

5 EFFICIENT IMPLEMENTATION OF THE SELECTION POLICIES

Both the maximum Sybil and view consistency policies offer ways to select a view, either the one claiming the most Sybils or the largest one that is γ_n -true, but brute-force examination of all $2^{|I|}$ receiver sets is infeasible. Instead, we describe $O(|I|^3)$ algorithms for both policies. In summary, both start by generating $O(|I|)$ candidate views (Algorithm 1). For the maximum Sybil policy, the one claiming the most Sybil identities is trivially identified. For the view consistency policy, Algorithm 2 is used to identify largest γ_n -consistent view.

Algorithm 1 Choose the receiver sets to consider

Require: i_0 is the identity running the procedure
Require: n is the desired receiver set size

- 1: $S \leftarrow \emptyset$
- 2: **for all** $i \in I$ **do**
- 3: $R \leftarrow \{i_0, i\}$
- 4: **for** $cnt = 3 \rightarrow n$ **do**
- 5: $R \leftarrow R \cup \{\text{RandElement}(V_{NS}(R))\}$
- 6: **end for**
- 7: $S \leftarrow S \cup \{R\}$
- 8: **end for**
- 9: **return** S ▷ with high probability, S contains a truthful receiver set

5.1 Candidate Receiver Set Selection

The only requirement for candidate receiver set selection is that at least one of the candidates must be truthful. Algorithm 1 selects $|I|$, size- n (we suggest $n = 4$) receiver sets of which at least one is truthful with high probability.




Fig. 6. Contours of probability that at least one of the receiver sets from Algorithm 1 is conforming.⁷ In the shaded areas, conditions required by the consistency policy or by Algorithm 1 are not met.

As illustrated in Figure 5, the algorithm starts with all $|I|$ size-2 receiver sets (lines 2–3) and builds each up to the full size- n by iteratively (line 4) adding a randomly selected identity from those indicated to be conforming at the prior lower dimensionality (line 5). At least $|C|$ of the initial size-2 receiver sets are conforming and after increasing to size- n , at least one is still conforming with high probability (graphed in Figure 6):

$$1 - \left(1 - \prod_{m=2}^{n-1} \frac{(1 - \gamma_m) \cdot |C| - (m - 1)}{|LNS| + (1 - \gamma_m) \cdot |C| - (m - 1)} \right)^{|C|}$$

In Figure 6 we use size-4 signalprints ($n = 4$) and $\gamma_4 = 0.05$, based on previous evaluation results [15], [17]. Shaded areas are ignored because we only consider $|C| > n$ and $|C| > |LNS|$. The former is a basic condition for Algorithm 1 to work, as there needs to be at least one size- n receiver set composed purely of conforming nodes. The latter is a basic condition of the view consistency policy (Condition 2). The contours are not smooth because $|C|$ and $|LNS|$ are discrete.

The signalprint threshold for this process is chosen to eliminate (nearly) all false negatives because the goal is to minimize the malicious-to-conforming ratio; false positives are harmless during the generation of candidate views. The complexity of a straightforward implementation is $O(|I|^3)$. Section 10 further discusses the runtime.

5.2 Finding the Largest γ_n -Consistent View

Given the $|I|$ candidate receiver sets, the next task is identifying the one generating a γ_n -true view, which,

⁷ For small $|C|$ and relatively large $|LNS|$ the probability can be increased by building $2 \cdot |I|$ or $3 \cdot |I|$ or more receiver sets instead.

Algorithm 2 Find receiver set permitting the largest γ_n -consistent subset

Require: S is the set of receiver sets generated by Algorithm 1

Require: $V_{\text{NS}}(R)$ for each $R \in \{\text{size-2 receiver sets}\}$ computed by Algorithm 1

Require: s is the initiator running the algorithm

```

1:  $(C, R_{\max}) \leftarrow (\infty, \text{null})$ 
2: for all  $R \in S$  do
3:   Compute RSSI ratio for each Sybil set in  $V_S(R)$ 
4:    $c \leftarrow 0$ 
5:   for all  $i \in V_{\text{NS}}(R)$  do
6:      $e \leftarrow 0$ 
7:      $n \leftarrow$  number of identities whose RSSI ratios reported by  $i$  do not match that for  $R$ 
8:     if  $\frac{|V_{\text{NS}}(R)|+n}{n} < \frac{1-2\gamma_n}{\gamma_n}$  then
9:        $e \leftarrow 1$ 
10:    end if
11:    if  $V(R)$  and  $V(\{i, s\})$  are not  $\gamma_2$ -similar then
12:       $e \leftarrow 1$ 
13:    end if
14:    if  $e = 1$  then
15:       $c \leftarrow c + 1$  ▷ exclude  $i$ 
16:    end if
17:  end for
18:  if  $c < C$  then
19:     $(C, R_{\max}) \leftarrow (c, R)$  ▷ new largest  $\gamma$ -consistent subset found
20:  end if
21: end for
22: return  $R_{\max}$ 

```

pursuant to Theorem 4, is that permitting the largest subset of I to be γ_n -consistent. Checking consistency by examining all $2^{|V_{\text{NS}}|}$ receiver sets is infeasible, so we make several observations leading to the $O(|I|^3)$ Algorithm 2. For each candidate receiver set (line 2), we determine how many identities must be excluded for the view to be γ_n -consistent (lines 3–17). The view excluding the fewest is γ_n -true and the desired classification (line 22).

The crux of the algorithm is lines 3–17, which use the following observations to efficiently determine which identities must be excluded.

- 1) Adding an identity to a receiver set can change the view in one direction only—an identity can go from Sybil to non-Sybil, but not vice versa—because uncorrelated RSSI vectors cannot become correlated by increasing the dimension.⁸
- 2) For identities a and b , $R \cup \{a\} \mapsto V(R)$ and $R \cup \{b\} \mapsto V(R)$ implies $R \cup \{a, b\} \mapsto V(R)$ because a and b must have the same RSSI ratios for the Sybils as R .

From these observations, we determine the excluded identities by computing, for each identity in $V_S(R)$, the RSSI ratio with an arbitrary sibling (line 3) and comparing against those reported by potential non-Sybils in $V_{\text{NS}}(R)$ (line 7). If the number not matching is too large (line 8), the view is not γ_n -consistent and the identity is excluded (line 15). It is also excluded if the receiver set consisting

8. This is not true for low dimension receiver sets severely affected by noise, but is for the size- $(n > 4)$ sets considered here.




Fig. 7. Distribution of RSSI variations in real-world deployment.

of just itself and the initiator is not γ_2 -similar to R (line 11).

5.3 Runtime in the Absence of Liars

In a typical situation with no liars, the consistency algorithm can detect the Sybils in $O(|I|^2)$ time. Since all identities are truthful, any chosen receiver set will be γ_n -consistent with no exclusions—clearly the largest possible—and thus the other $|I| - 1$ also-truthful receiver sets need not be checked. With lying attackers present, the overall runtime is $O(|I|^3)$, as each algorithm takes $O(|I|^3)$ time.

6 CLASSIFICATION PERFORMANCE AGAINST OPTIMAL ATTACKERS

Both view selection policies depend directly on the unpredictability of RSSIs, because collapsing identities requires knowing the observations of the initiator, as explained in Section 4.1. An intelligent attacker can attempt educated guesses, resulting in some successful collapses. In this section, we evaluate the two selection policies against optimal attackers, as defined in Sections 6.2 and 6.3.

6.1 RSSI Unpredictability

Accurately guessing RSSIs is difficult because the wireless channel varies significantly with small displacements in location and orientation (*spatial variation*) and environmental changes over time (*temporal variation*) [13], [26]. Pre-characterization could account for spatial variation, but would be prohibitively expensive at the needed spatial and orientation granularity (6 cm [27] and 3° for our test devices).

We empirically determined the RSSI variation for human-carried smartphones by deploying experimental phones to eleven graduate students in two adjacent offices and measuring pairwise RSSIs for fifteen hours. The observed distribution of deviations⁹, shown in Figure 7, is roughly normal with a standard deviation of 7.3 dBm,

9. For each pair of transceivers, we subtracted the mean of all their measurements to get the deviations and took the distribution of the pairwise deviations.

in line with other real-world measurements for spatial and orientation variations (4–12 dBm and 5.3 dBm [13]). We use this distribution to model the attacker uncertainty of RSSIs, corresponding to an attacker who accumulates knowledge of pairwise RSSIs by observing values reported in past tests.

6.2 Optimal Attacker Strategy—Maximum Sybil Policy

Theorem 1 shows that the performance of the maximum Sybil policy is inversely related to the number of collapsed non-Sybil identities. Therefore, the optimal attacker tries to collapse as many as possible. We give two observations about this goal.

- 1) More distinct guesses increase the probability of success, so an optimal attacker partitions its (mostly Sybil) identities, with each group making a different guess.
- 2) Smaller group size increases the number of groups, but decreases the probability that the group is considered—recall that Algorithm 1 generates only $|I|$ of the possible $2^{|I|}$ candidate receiver sets.

Consequently, there is an optimal group size that maximizes the total number of groups (guesses) produced by Algorithm 1, which we obtained via Monte Carlo simulations. We model the initiator’s RSSI observation as a random vector whose elements are drawn i.i.d. from the Gaussian distribution in Figure 7. Given the total number of guesses, the best choices are the vectors with the highest joint probabilities. The performance against this strategy is discussed in Section 6.4.

6.3 Optimal Attacker Strategy—View Consistency Policy

The view consistency policy depends on Condition 3 holding, i.e., all consistent views must correctly classify at least $|LNS| + 1$ conforming identities. In this section we quantify the probability that it holds against an optimal attacker. To break Condition 3, an attacker must generate a consistent view that collapses at least $|C| - |LNS|$ conforming identities. We give three observations about the optimal attacker strategy for this goal.

- 1) Collapsing $|C| - |LNS|$ identities is easiest with larger $|LNS|$. Thus, the optimal attacker uses only one physical node to claim Sybils—the others just lie.
- 2) For a particular false view to be consistent, all supposedly non-Sybil identities must indict the same identities, e.g., have the same RSSI guesses for the collapsed conforming identities. The optimal attacker must divide its (mostly Sybil) identities into groups, each using a different set of guesses.
- 3) More groups increases the probability of success, but decreases the number of Sybils actually accepted, as each group is smaller.




Fig. 8. Contours of a lower bound on the probability that Condition 3 holds under an optimal attacker strategy with the attacker’s knowledge of RSSIs modeled as a normal distribution with standard deviation 7.3 dBm.

We assume the optimal attacker wishes to maximize the probability of success and thus uses minimum-sized groups (three identities, for size-4 signalprints).

For each group, the attacker must guess RSSI values for the conforming identities with the goal of collapsing at least $s \triangleq |C| - |LNS|$ of them. There are $\binom{|C|}{s}$ such sets and the optimal attacker guesses values that maximize the probability of at least one (across all groups) being correct. The first group is easy; the $|C|$ guesses are simply the most likely values, i.e., the expected values for the conforming identities’ RSSIs, under the uncertainty distribution.

For the next (and subsequent) groups, the optimal attacker should pick the next most likely RSSI values for each of the $\binom{|C|}{s}$ sets. However, the sets share elements (only $|C|$ RSSIs are actually guessed), so the attacker must determine the most probable sets that are compatible.

This is a non-trivial problem but it is conceivable that an attacker might be able to solve it. In order to model the strongest possible attack, we assume that all sets are compatible. For example, one group can guess -78 dBm and -49 dBm RSSIs for nodes a and b , but -82 dBm and -54 dBm RSSIs for nodes a and c . This is impossible because it would imply that the group simultaneously guess both -78 dBm and -82 dBm for node a . Recall that, each group only have the opportunity to guess the RSSI for a node once during the Mason test. Therefore, the assumption that all sets are compatible allows us to model an attack that is at least as strong as, but likely stronger than, any realizable attack. This leads to a conservative lower bound on the probability that the attacker fails—any feasible, optimal strategy is less likely to succeed.

Figure 8 shows contours of this lower bound on the probability that Condition 3 holds as a function of $|C|$ and $|LNS|$, obtained via Monte Carlo simulations. The




Fig. 9. The final Sybil ratio, i.e., fraction of accepted identities that are Sybil, produced by the maximum Sybil policy against an optimal attacker strategy.

initiator’s RSSI observation is modeled as a random vector whose elements are drawn i.i.d. from the Gaussian distribution in Figure 7. The $|C| \leq |LNS|$ region is shaded, because the view consistency policy fails there (recall Condition 2). When the conforming nodes outnumber the attacker nodes by at least $1.5\times$ —the expected case in real networks—the condition holds with very high probability. In practice, it will hold with even higher probability, as this is a lower bound.

6.4 Performance Comparison of Both Policies

We use Monte Carlo simulations to compare the performance of the two policies against the optimal attackers, quantified as the *final Sybil ratio*, the fraction of accepted identities that are Sybil. We model the attacker’s knowledge of the initiator’s RSSIs as a random vector whose elements are drawn i.i.d. from the Gaussian distribution in Figure 7, which conservatively assumes fine-grained temporal and spatial characterization (see Section 6.1). We expect real-world attackers to have less knowledge, leading to even better classification performance.

Our procedure for generating candidate receiver sets (Algorithm 1) works best when conforming nodes outnumber physical attackers. This condition should normally hold in real-world networks (it is the major motivation for a Sybil attack), so for both policies, we report results assuming that it does.

Figure 9 graphs the final Sybil ratio of the maximum Sybil policy, which roughly corresponds to the ratio of collapsed conforming nodes ($\frac{|V_S \cap NS|}{|C|}$). The performance does not depend on the number of physical attackers. The Sybil ratio decreases to 0.05–0.2 when $|C| > 10$. When $|C| < 10$, the Sybil ratio is high (0.2–0.5), despite elimination of most Sybil identities (92%–99%). This behavior is due to the ease of guessing low-dimension random vectors.

Figure 10 shows the final Sybil ratio of the consistency policy. Again, the $|C| \leq |LNS|$ region is shaded as the policy simply fails in this case. Performance increases rapidly with the ratio of conforming nodes to physical attackers—recall the attacker needs to collapse $|C| - |LNS|$




Fig. 10. Contours showing the final Sybil ratio for the view consistency policy against an optimal attacker strategy. The dashed line corresponds to situations where this policy has the same performance as the maximum Sybil policy.

identities to break Condition 3. For example, the final Sybil ratio drops below 10^{-6} when $\frac{|C|}{|LNS|+1} \geq 1.6$. As the collapse rate is usually below 0.2 (see Figure 9 when $|C| > 10$), we observe good performance when $|C| - |LNS| \geq 0.2|C|$ (below the 0.05 contour). The dashed line (roughly $\frac{|C|}{|LNS|+1} = 1.2$) indicates the situations where both policies perform equally. Below it, the consistency policy performs better than the maximum Sybil policy and above it does worse.

The view consistency policy is superior when conforming nodes are expected to outnumber attacker nodes by at least $1.2\times$, the common case in urban environments. The maximum Sybil policy remains viable when the number of physical attackers is comparable to (or even larger than) that of the conforming nodes. We suggest users of the Mason test consider their application knowledge when choosing a policy.

7 DETECTING MOVING ATTACKERS

A mobile attacker can defeat signalprint comparison by changing locations or orientations between transmissions to associate distinct signalprints with each Sybil identity. Instead of restricting the attack model to only stationary devices, we protect against moving attacks by detecting moving nodes. Moving nodes are treated as non-conforming, in essence, and will not be able to participate in network protocols, until stationary enough to be tested for Sybilness again. Fortunately, in the networks we consider, most conforming nodes (e.g., human-carried smartphones and laptops) are stationary over most short time-spans (1–2 min), due to human mobility habits. Thus, multiple transmissions should have the same signalprints[15]. From this observation, we develop a protocol to detect moving attackers.




Fig. 11. Contours showing the response time (in ms, 99th percentile) to precisely switch between two positions required to defeat the challenge-response moving node detection.

Again, the lack of trusted observations is troublesome. Instead of comparing signalprints, we compare the initiator’s observations: all transmissions from a conforming node should have the same RSSI. As shown in Section 9, this simple criterion yields acceptable detection.

The protocol collection phase (Figure 2a) is extended to request multiple probe packets (e.g., 14) from each identity in a pseudo-random order (see Section 8.1). During the classification phase (Figure 2c) each participant rejects any identity with a large RSSI variation across its transmissions (specifically, a standard deviation larger than 2.5 dBm). In essence, an attacker is challenged to quickly and precisely switch between the multiple positions associated with its Sybil identities (6 cm location precision according to coherence length theory [27] and 3° orientation precision according to our measurements).

Figure 11 plots the required response time for an attacker to pass the challenge. Random sequences of probe requests are generated via Monte Carlo simulations and the required response time is calculated accordingly. Given human reaction times [28], reliably mounting such an attack would require specialized hardware—precise electromechanical control or beam steering antenna arrays—that is outside our attack model and substantially more expensive to deploy than compromised commodity devices.

8 THE MASON TEST

This section describes the full Mason test protocol, an implementation of the concepts introduced in the previous sections. There are four main requirements on the protocol.

- 1) Conforming neighbors must be able to participate. That is, selective jamming of conforming identities

must be detectable.

- 2) Probe packets must be transmitted in pseudo-random order. Further, each participant must be able to verify that no group of identities controlled the order.
- 3) Moving identities must be rejected. To save energy and time, conforming nodes that are moving when the protocol begins should not participate.
- 4) Attackers must not know the RSSI observations of conforming identities when constructing lies.

We assume a known upper bound on the number of conforming neighbors, i.e., those within the one-hop transmission range. In most applications, a bound in the hundreds (we use 400 in our experiments) will be acceptable. If more identities attempt to participate, the protocol aborts and no classification is made. This appears to open a denial-of-service attack. However, we do not attempt to prevent, instead only detect, DOS attacks, because one such attack—simply jamming the wireless channel—is unpreventable (with commodity hardware). See Section 10 for more discussion.

The rest of this section describes the two components of the protocol: collection of RSSI observations and Sybil classification. We assume one identity, the initiator, starts the protocol and leads the collection, but all identities still individually and safely perform Sybil classification.

8.1 Collection of RSSI Observations

Phase I: Identity Collection. The first phase gathers participating neighbors, ensuring that no conforming identities are jammed by attackers. The initiator sends a REQUEST message stating its identity, e.g., a public key. All stationary neighbors respond with their identities via HELLO-I messages, ACKed by the initiator. Unacknowledged HELLO-Is are re-transmitted. The process terminates when the channel is idle—indicating all HELLO-I’s were received and ACKed. If the channel does not go idle before a timeout (e.g., 15 seconds), the protocol aborts because an attacker may be selectively jamming some HELLO-Is. The protocol also aborts if too many identities join, e.g., 400.

Phase II: Randomized Broadcast Request: The second phase is the challenge-response protocol to collect RSSI observations for motion detection and Sybil classification. First, each identity contributes a (difficult to predict) random value¹⁰; all are hashed together to produce a seed to generate the random sequence of broadcast requests issued by the initiator. Specifically, it sends a TRANSMIT message to each participant in the random sequence, who must quickly broadcast a signed HELLO-II, e.g., within 10 ms in our implementation¹¹. Each participant

10. Even if attackers do not comply, conforming participants can verify that their own random submissions resulted in a random sequence and therefore trust the test results.

11. 10 ms is larger than the typical roundtrip time for 802.11b with packets handled in interrupt context for low-latency responses. These packets can be signed ahead of time and cached—signatures do not need to be computed in the 10 ms interval.

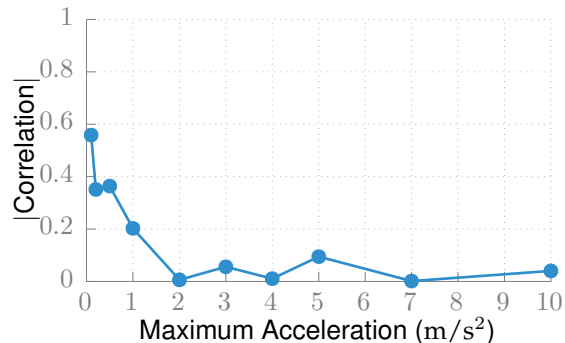


Fig. 12. RSSI correlation as a function of the maximum device acceleration between observations.

records the RSSIs of the HELLO-II messages it hears. Some identities will not hear each other; this is acceptable because the initiator needs observations from only three other conforming identities. $|I| \times s$ requests are issued, where s is large enough to ensure a short minimum duration between consecutive requests for any two pairs of nodes, e.g., 14 in our tests. An identity that fails to respond in time might be an attacker attempting to change physical position and is rejected.

Phase III: RSSI Observations Report. In the third phase, the RSSI observations are shared. First, each identity broadcasts a hash of its observations. Then the actual values are shared. Those not matching the respective hash are rejected, preventing attackers from using the reported values to fabricate plausible observations. The same mechanism from Phase 1 is used to detect selective jamming.

8.2 Sybil Classification

Each participant performs Sybil classification individually. First, the identity verifies that its observations were not potentially predictable from those reported in prior rounds, possibly violating Condition 3. Correlation in RSSI values between observations decreases with motion between observations, as shown by our experiments (Figure 12). Thus, a node only performs Sybil classification if it has strong evidence that the current observations are uncorrelated with prior ones¹², i.e., it has observed an acceleration of at least 2 m/s^2 .

Classification is a simple application of the methods of Section 7 and Section 5. Each identity with an RSSI variance across its multiple broadcasts higher than a threshold is rejected. Then, Algorithm 1 and Algorithm 2 are used to identify a γ -true Sybil classification over the remaining, stationary identities.

9 PROTOTYPE AND EVALUATION

We implemented the Mason test as a Linux kernel module and tested its performance on HTC Magic Android

12. Note that although we did not encounter this case in our experiments, it is conceivable that some devices will return to the same location and orientation after motion.

TABLE 2
Thresholds for Signalprint Comparison and Motion Filtering

Name	Threshold (dBm)
Signalprint Distance	dimension-2 0.85 dimension-3 3.6 dimension-4 1.2
RSSI Standard Deviation	2.5

smartphones in various operating environments. It sits directly above the 802.11 link layer, responding to requests in interrupt context, to minimize response latency for the REQUEST-HELLO-II sequence (12 ms roundtrip time on our hardware). The classification algorithms are implemented in Python. Unlike the described protocol, mobile conforming nodes participated in all tests (i.e., nodes did not monitor their own motion and decline to participate when moving), giving us data to tune the motion filter and characterize the impact of node motion on the classifier performance.

The goal of this section is to evaluate the overall performance of our system in normal settings, which is mainly dependent on the wireless environment. We therefore evaluated the Mason test in four different environments.

Office I Eleven participants in two adjacent offices for fifteen hours.

Office II Eleven participants in two adjacent offices in a different building for one hour, to determine whether performance varies across similar, but non-identical environments.

Cafeteria Eleven participants in a crowded cafeteria during lunch. This was a rapidly-changing wireless environment due to frequent motion of the cafeteria patrons.

Outdoor Eleven participants meeting in a cold, open, grassy courtyard for one hour, capturing the outdoor environment. Participants moved frequently to stay warm.

In each environment, we conducted multiple trials with one Sybil attacker¹³ generating 4, 20, 40, and 160 Sybil identities. The ratio of conforming to attacking nodes is held constant, as it does not affect performance (assuming at least one true view is generated by Algorithm 1). The gathered traces were split into testing and training sets.

Note that we do not study the performance under collapsing attacks here, as it depends on the number of conforming and attacking nodes. We have too few experimental devices to meaningfully vary those counts. Instead, we use Monte Carlo simulations in Section 6 to independently evaluate performance against such attacks.

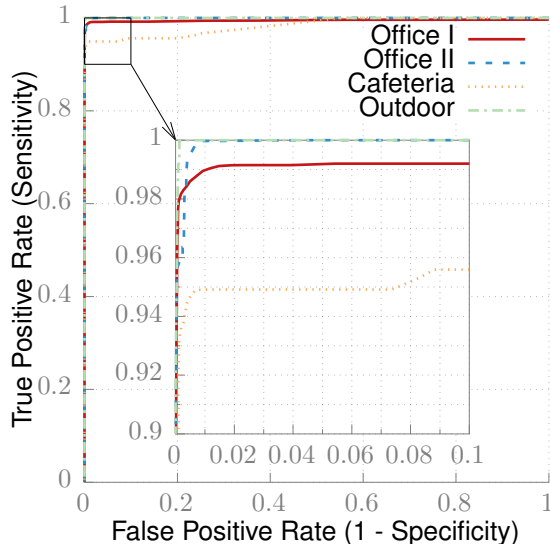


Fig. 13. ROC curve showing the classification performance of signalprint comparison in different environments for varying distance thresholds. Only identities that passed the motion filter are considered. The knees of the curves all correspond to the same thresholds, suggesting that the same value can be used in all locations.

TABLE 3
Classification Performance

Environment	Sensitivity (%)	Specificity (%)
Office I	99.6	96.5
Office II	100.0	87.7
Cafeteria	91.4	86.6
Outdoor	95.9	61.1

9.1 Selection and Robustness of Thresholds

The training data were used to determine good motion filter and signalprint distance thresholds, shown in Table 2.

The motion filter threshold was chosen such that at least 95% of the conforming participants (averaged over all training rounds) in the low-motion Office I environment would pass. This policy ensures that conforming smartphones, which are usually left mostly stationary, e.g., on desks, in purses, or in the pockets of seated people, will usually pass the test. Devices exhibiting more motion (i.e., a standard deviation of RSSIs at the initiator larger than 2.5 dBm)—as would be expected from an attacker trying to defeat signalprint detection—will be rejected.

The signalprint distance thresholds were chosen by evaluating the signalprint classification performance at various possible values. Figure 13 shows the ROC curves for size-4 receiver sets (a “positive” is an identity classified as Sybil). Note that the true positive and false positive rates consider only identities that passed the motion

filter, in order to isolate the effects of the signalprint distance threshold. The curves show that a good threshold has performance in line with prior work [15], [17], as expected.

In all environments, the knees of the curve correspond to the same thresholds, suggesting that these values can be used in general, across environments. A possible explanation is that despite environment differences, the signalprint distance distributions for stationary Sybil siblings are identical. All results in this paper use these same thresholds, shown in Table 2.

9.2 Classification Performance

The performance of the full Mason test—motion filtering and signalprint comparison—is detailed by the confusion matrices in Figure 14. Note that we count all rejected identities, including both Sybil and moving identities, as Sybil. Many tests were conducted in each environment, so average percentages are shown instead of absolute counts. To evaluate the performance, we consider two standard classification metrics derived from these matrices, *sensitivity* (percentage of Sybil identities correctly identified) and *specificity* (percentage of conforming identities correctly identified).

Note that 100% sensitivity is not necessary. Most protocols that would use Mason require a majority of the participants to be conforming. The total number of identities is limited (e.g., to 400), so rejecting most of the Sybils and accepting most of the conforming identities is sufficient to meet this requirement.

Table 3 shows the performance for all four environments. The Mason test performs best in the low-motion indoor environments, with over 99.5% sensitivity and over 85% specificity. The sensitivity in the cafeteria environment is just 91.4%, likely due to the rapid and frequent changes in the wireless environment resulting from the motion of cafeteria patrons. In the outdoor environment, with all participants (including attackers) moving, the sensitivity is 95.9%, and the specificity is 61.1% with all the false rejections caused by motion.

The outdoor experiment is an extreme case where we pay the cost of rejecting moving conforming nodes to defeat motion attacks. The result is acceptable because our goal is to produce a set of non-Sybil identities to be used safely by other protocols: accepting a swarm of moving Sybil identities is much worse than temporarily rejecting some conforming nodes that are currently moving.

An identity is classified as Sybil for three reasons: it has similar signalprints to another, the initiator has too few RSSI reports to form a signalprint, or it is in motion. Figure 15 shows the relative prevalence of these three causes for falsely rejecting conforming nodes. Not surprisingly, the first cause—collapsing—is rare, occurring only in the first office environment. Missing RSSI reports is an issue only in the environments with significant obstructions (the indoor offices) and accounts for about half of these false rejections. In the open

13. As discussed in Section 4 and Section 6, additional physical nodes are best used as lying, non-Sybils.




Fig. 14. Confusion matrices detailing the classifier performance in the four environments. S is Sybil and C is conforming. Multiple tests were run in each environment, so mean percentages are shown instead of absolute counts.




Fig. 15. Relative frequencies of the three causes of false positives.

cafeteria and outdoor environments, all false rejections are due to participant motion.

9.3 Overhead Evaluation

Figures 16a and 16b show the runtime and energy overhead for the Mason test collection phase, with the stacked bars separating the costs by sub-phase. The protocol runs infrequently (once every hour is often sufficient), so runtimes of 10–90 seconds are acceptable. Likewise, smartphone energy consumption is acceptable, with the extreme 18J consumption for 400 identities representing 0.01% of the 17.500J capacity of a typical smartphone batteries.

Figure 16c show the classification phase overheads for 2–100 identities. Classification consumes much less energy than collection, so its overhead is also acceptable. For more than 100 participants, costs become excessive due to the $O(n^3)$ scaling behavior¹⁴. Limiting participation to 100 identities may be necessary for energy-constrained devices, but will generally not reduce performance because having 100 non-Sybil, one-hop neighbors is rare.

The periodic accelerometer sampling used to measure motion between Mason test rounds consumes 5.2% of battery capacity in an 18 h period of use before recharging.

14. A native-C implementation might scale to 300–400 identities.

10 DISCUSSION


Sybil classification from untrusted observations is difficult and the Mason test is not a silver bullet. Not requiring trusted observations is a significant improvement, but the test’s limitations must be carefully considered before deployment. As with any system intended for real-world use, some decisions try to balance system complexity and potential security weaknesses. In this section, we discuss these trade-offs, limitations, and related concerns.

High Computation Time: The collection phase time is governed by the 802.11b-induced 12 ms per packet latency and the classification runtime grows quickly with the number of identities, $O(|I|^3)$. Although typically fast (e.g., <5 s for 5–10 nodes), the Mason test is slower in high density areas (e.g., 40 s for 100 nodes). However, it should be run infrequently, e.g., once or twice per hour. Topologies change slowly (most people change locations infrequently) and many protocols requiring Sybil resistance can handle the lag—they need only know a subset of the current non-Sybil neighbors.

Easy Denial-of-Service Attack: An attacker can force the protocol to abort by creating many identities or jamming transmissions from the conforming identities. We cannot on commodity 802.11 devices solve another denial-of-service attack—simply jamming the channel—so defending against these more-complicated variants is ultimately useless. Most locations will at most times be free of such attackers—the Mason test provides a way to verify this condition, reject any Sybils, and let other protocols operate knowing they are Sybil-free.

Requires Several Conforming Neighbors: The Mason test requires true RSSI observations from some neighbors (i.e., 3) and is easily defeated otherwise. Although a detailed treatment is beyond the scope of this paper, we do note that protocols incorporating the Mason test can mitigate this risk by (a) a priori estimation of the distribution of the number of conforming neighbors and (b) careful composition of results from multiple rounds to bound the failure probability.

Limit On Total Identities: This limit (e.g., 400) is unfortunately necessary to detect when conforming nodes are being selectively jammed while keeping the test duration short enough that most conforming nodes remain stationary. We believe that most wireless networks



(a) Runtime in seconds of the collection phase.

(b) Energy consumption in joules of the collection phase.

(c) Runtime and energy consumption of the classification phase.

Fig. 16. Overhead of the collection and classification phases. The stacked bars partition the cost among the participant collection (HELLO I), RSSI measurement (HELLO II), and RSSI observation exchange (RSST) steps.

have typical node degrees well below 400.

Messages Must Be Signed: Packets sent during the collection phase are signed, which can be very slow with public key schemes. However, this is easily mitigated by (a) pre-signing the packets to keep the delay off the critical path or (b) using faster secret-key-based schemes.

Pre-Characterization Reveals RSSIs: An attacker could theoretically improve its collapsing probability by pre-characterizing the wireless environment. We believe such attacks are impractical because the required spatial granularity is about 6 cm, the device orientation is still unknown, and environmental changes (e.g., people moving) reduces the usefulness of prior characterization.

Prior Rounds Reveal RSSI Information: The protocol defends against this. Conforming nodes do not perform classification if their RSSI observations are correlated with the prior rounds (see Section 8.2).

High False Positive Rates: With the motion filter, the false positive rate can be high, e.g., 20% of conforming identities rejected in some environments. We believe this is acceptable because most protocols requiring Sybil resistance need only a subset of honest identities. For example, if for reliability some data is to be spread among multiple neighbors, it is acceptable to spread it among a subset chosen from 80%, rather than all, of the non-Sybils.

11 CONCLUSION

We have described a method to use signalprints to detect Sybil attacks in open ad hoc and delay-tolerant networks without requiring trust in any other node or authority. We use the inherent difficulty of predicting RSSIs to separate true and false RSSI observations reported by one-hop neighbors. Attackers using motion to defeat the signalprint technique are detected by requiring low-latency retransmissions from the same position.

The Mason test was implemented on HTC Magic smartphones and tested with human participants in three environments. It eliminates 99.6%–100% of Sybil identities in office environments, 91% in a crowded high-motion cafeteria, and 96% in a high-motion open outdoor environment. It accepts 88%–97% of conforming identities in the office environments, 87% in the cafeteria, and 61%

in the outdoor environment. The vast majority of rejected conforming identities were eliminated due to motion.

REFERENCES

- [1] P. Hui, J. Crowcroft, and E. Yoneki, "BUBBLE rap: Social-based forwarding in delay tolerant networks," *IEEE Trans. Mobile Computing*, vol. 10, no. 11, pp. 1576–1589, Nov. 2011.
- [2] Y. Xiang, et al., "Collaborative calibration and sensor placement for mobile sensor networks," in *Proc. Int. Conf. Information Processing in Sensor Networks*, Apr. 2012, pp. 73–84.
- [3] P. Gardner-Stephen, "The Serval project: Practical wireless ad-hoc mobile telecommunications," Flinders University, Adelaide, South Australia, Tech. Rep., Aug. 2011.
- [4] J. Douceur, "The Sybil attack," in *Proc. Int. Wkshp. Peer-to-Peer Systems*, Mar. 2002, pp. 251–260.
- [5] J. Newsome, et al., "The Sybil attack in sensor networks: Analysis & defenses," in *Proc. Int. Conf. Information Processing in Sensor Networks*, Apr. 2004, pp. 259–268.
- [6] B. N. Levine, C. Shields, and N. B. Margolin, "A survey of solutions to the Sybil attack," Department of Computer Science, University of Massachusetts Amherst, Amherst, MA, Tech. Rep., Oct. 2006.
- [7] H. Zhou, M. Mutka, and L. Ni, "Multiple-key cryptography-based distributed certificate authority in mobile ad-hoc networks," in *Proc. Global Telecommunications Conf.*, Nov. 2005.
- [8] M. Ramkumar and N. Memon, "An efficient key predistribution scheme for ad hoc network security," *IEEE J. Selected Areas in Communications*, vol. 23, pp. 611–621, Mar. 2005.
- [9] N. Borisov, "Computational puzzles as Sybil defenses," in *Proc. Int. Conf. Peer-to-Peer Computing*, Sept. 2006, pp. 171–176.
- [10] F. Li, et al., "SybilControl: Practical Sybil defense with computational puzzles," in *Proc. Wkshp. Scalable Trusted Computing*, Oct. 2012.
- [11] H. Yu, et al., "SybilGuard: defending against Sybil attacks via social networks," in *Proc. ACM SIGCOMM Computer Communication Review*, Sept. 2006, pp. 267–278.
- [12] H. Yu, et al., "SybilLimit: A near-optimal social network defense against Sybil attacks," in *Proc. Symp. Security and Privacy*, May 2008, pp. 3–17.
- [13] T. S. Rappaport, *Wireless Communications: Principles & Practice*. Prentice-Hall, NJ, 2002.
- [14] A. Haeberlen, et al., "Practical robust localization over large-scale 802.11 wireless networks," in *Proc. Int. Conf. Mobile Computing and Networking*, Sept. 2004, pp. 70–84.
- [15] L. Xiao, et al., "Channel-based detection of Sybil attacks in wireless networks," *IEEE Trans. Information Forensics and Security*, vol. 4, no. 3, pp. 492–503, Sept. 2009.
- [16] D. B. Faria and D. R. Cheriton, "Detecting identity-based attacks in wireless networks using signalprints," in *Proc. Wkshp. Wireless Security*, Sept. 2006, pp. 43–52.
- [17] M. Demirbas and Y. Song, "An RSSI-based scheme for Sybil attack detection in wireless sensor networks," in *Proc. Int. Symp. on a World of Wireless, Mobile, and Multimedia*, June 2006, pp. 564–570.
- [18] Z. Li, et al., "Securing wireless systems via lower layer enforcements," in *Proc. Wkshp. Wireless Security*, Sept. 2006, pp. 33–42.

- [19] Q. Li and W. Trappe, "Detecting spoofing and anomalous traffic in wireless networks via forge-resistant relationships," *IEEE Trans. Information Forensics and Security*, vol. 2, no. 4, pp. 793–803, Dec. 2007.
- [20] Y. Chen, et al., "Detecting and localizing identity-based attacks in wireless and sensor networks," *IEEE Trans. Vehicular Technology*, vol. 5, no. 5, pp. 2418–2434, June 2010.
- [21] T. Suen and A. Yasinsac, "Peer identification in wireless and sensor networks using signal properties," in *Proc. Int. Conf. Mobile Adhoc and Sensor Systems*, Nov. 2005, pp. 826–833.
- [22] S. Lv, et al., "Detecting the Sybil attack cooperatively in wireless sensor networks," in *Proc. Int. Conf. Computational Intelligence and Security*, Dec. 2008, pp. 442–446.
- [23] M. S. Bouassida, et al., "Sybil nodes detection based on received strength variations within VANET," *Int. J. Network Security*, vol. 9, no. 1, pp. 22–33, July 2009.
- [24] D. Gesbert, et al., "From theory to practice: An overview of MIMO space-time coded wireless systems," *IEEE J. Selected Areas in Communications*, vol. 21, no. 3, pp. 281–302, Apr. 2003.
- [25] Y. Liu, D. R. Bild, and R. P. Dick, "Extending channel comparison based Sybil detection to MIMO systems," Dept. of Electrical Engineering and Computer Science, University of Michigan, Tech. Rep., <http://www.davidbild.org/publications/liu13dec.pdf>.
- [26] H. Hashemi, D. Lee, and D. Ehman, "Statistical modeling of the indoor radio propagation channel – part II," in *Proc. Vehicular Technology Conf.*, May 1992, pp. 839–843.
- [27] T. S. Rappaport, S. Y. Seidel, and K. Takamizawa, "Statistical channel impulse response models for factory and open plan building radio communication system design," *IEEE Trans. on Communications*, vol. 39, no. 5, pp. 794–806, May 1991.
- [28] "Reaction time statistics," <http://www.humanbenchmark.com/tests/reactiontime/stats.php>.

Yue Liu received the BE and MS degrees in electrical engineering from the University of Electronic Science and Technology of China and the

University of Michigan, in 2007 and 2011, respectively. She is currently working towards a PhD degree in electrical engineering: systems in the University of Michigan. Her research interest includes wireless networks, delay-tolerant and mobile ad hoc networks, and network security.

David R. Bild (S'08-M'15) received the BS and MS degrees in computer engineering in 2007 and 2008 from Northwestern University and the PhD degree in computer science and engineering in 2014 from the University of Michigan. He has published on a variety of topics including integrated circuit reliability, digital systems testing, secure communication networks, and social network analysis.

Robert P. Dick (S'95-M'02) is an Associate Professor of EECS at the University of Michigan. His research focuses on embedded systems. He received his Ph.D. degree from Princeton University and his B.S. degree from Clarkson University. He was a Visiting Professor of EE at Tsinghua University, a Visiting Researcher at NEC Labs America, and an Associate Professor of EECS at Northwestern University. Robert received an NSF CAREER award and won a departmental Teacher of the Year award. His technology won a Computerworld Horizon Award and two of his publications won best paper awards.

Z. Morley Mao is an Associate Professor in the Department of Electrical Engineering and Computer Science at University of Michigan. She received her B.S., M.S., and Ph.D degree all from University of California at Berkeley. She is a recipient of the NSF CAREER Award, Sloan Fellowship, and the IBM Faculty Partnership Award. Her research interests include networking, systems, and security.

Dan S. Wallach is a Professor in the Department of Computer Science and a Rice Scholar in the Baker Institute for Public Policy at Rice University. He received his B.S. at the University of California at Berkeley and his M.A. and Ph.D at Princeton University. He is a recipient of the NSF CAREER Award and a variety of industrial partnership awards. His research interests consider computer systems security.