Detecting the Sybil Attack in Mobile Ad hoc Networks

Chris Piro
Dept. of Computer Science
Georgetown Univ.
cjp23@georgetown.edu

Clay Shields
Dept. of Computer Science
Georgetown Univ.
clay@georgetown.edu

Brian Neil Levine
Dept. of Computer Science
Univ. Massachusetts Amherst
brian@cs.umass.edu

Abstract

Mobility is often a problem for providing security services in ad hoc networks. In this paper, we show that mobility can be used to enhance security. Specifically, we show that nodes that passively monitor traffic in the network can detect a Sybil attacker that uses a number of network identities simultaneously. We show through simulation that this detection can be done by a single node, or that multiple trusted nodes can join to improve the accuracy of detection. We then show that although the detection mechanism will falsely identify groups of nodes traveling together as a Sybil attacker, we can extend the protocol to monitor collisions at the MAC level to differentiate between a single attacker spoofing many addresses and a group of nodes traveling in close proximity.

1 Introduction

NUMEROUS protocols exist for forming *ad hoc networks* among cooperative mobile, radio-equipped nodes [25, 14, 13]. Many ad hoc routing protocols have been secured using reputation schemes [3] or threshold security schemes [32, 16, 15] that rely on there being a limited number of attackers in the group and that assume each radio represents a different individual. However, the broadcast nature of radio allows a single node to pretend to be many nodes simultaneously by using many different addresses while transmitting.

This attack, an example [22] of what is called the Sybil attack [10], can easily defeat reputation [9] and threshold [10] protocols intended to protect against it. Douceur has shown that there is no practical defense against the attack; even a central authority (such as a PKI) must ensure that each *identity* is actually one *entity* — this requires costly manual intervention, which restricts the number of identities that can be managed. In contrast, protocols for detection do not suffer from such limitations. Moreover, detection is complementary to any method that attempts protection.

In this paper, we show that the mobility of nodes in a wireless network can be used to detect and identify nodes that are part of a Sybil attack. We rely on the fact that while individual nodes are free to move independently, all identities of a single Sybil attacker are bound to a single physical node and must move together. We propose two initial methods, both passive, that can be run on standard, inexpensive equipment without any special antennae or hardware and with only very loose clock synchronization.

In the first method, called *Passive Ad hoc Sybil Identity Detection* (PASID), a single node can detect Sybil attacks by recording the identities, namely the MAC or IP addresses of other nodes it hears transmitting. Over time, the node builds a profile of which nodes are heard together, which helps reveal Sybil attackers. We show through simulation that in networks with sufficient connectivity and mobility PASID can produce close to 100% accuracy in identifying the various attacker identities while avoiding any false positives. As the network becomes more dense, with more nodes in less space, the false positive rate increases; as it becomes more sparse, the accuracy rate declines as each node has fewer chances to hear its neighbors. To combat this, we show that multiple trusted nodes can share their observations to increase the accuracy of detection over a shorter time or in a more-sparsely connected network.

Our second method, *PASID with Group Detection* (PASID-GD), extends our approach and reduces false positives that can occur when a group of nodes moving together is falsely identified as a single Sybil attacker. By monitoring collisions at the MAC level we show that we can differentiate these cases. This approach is successful because an attacker operating over a single channel can transmit only serially, whereas independent nodes can transmit in parallel, creating detectably higher collision rates.

Our work follows an increasing body of work that sees mobility itself as an opportunity for improving the performance of mobile networking [8, 6].

2 Related Work

The Sybil attack can occur in a distributed system that operates without a central authority to verify the identities of each communicating entity [10]. Because each entity is only

This work was supported in part by NSF grants CNS-0133055, CNS-0534618, and CNS-0087639.

aware of others through messages over a communication channel, a Sybil attacker can assume many different identities by sending messages with different identifiers.

An entity in the system can attempt to determine if some set of entities are distinct by testing their resource limits, but this is problematic. If a single Sybil attacker pretends to be multiple entities, it may not have the same computational, storage, and bandwidth capabilities as multiple independent entities. However, testing based on such an assumption requires an accurate model of the attacker's resources. A Sybil attacker that has more resources than expected can impersonate a number of entities proportional to the amount its resources are underestimated. Similarly, a set of entities that are more resource-constrained than expected may fail to prove their independence.

The testing entity might also attempt to verify identity and independence indirectly by asking entities to vouch for each other. This strategy is prone to the Sybil attack because multiple entities can be the multiple identities of one or more Sybil attackers.

Newsome, et al [22] proposed several methods for detecting Sybil entities in a sensor network. They present an excellent discussion of the threat the Sybil attack poses to sensor networks, all of which apply to routing for ad hoc mobile networks. In contrast to the methods we propose, the detection techniques they proposed are *active* tests that require the participation of the neighboring nodes by asking them to respond to queries on assigned channels or to carry pre-distributed keys. Such query/response resource tests are a challenge to undertake in a mobile environment where neighbors legitimately may change with great frequency and without notice. Pre-distributing keys in an ad hoc network may not be possible if the nodes to not originate from the same source or are not all present for a key initialization phase. Regardless, a reliance on keys to detect or prevent a Sybil attack is based on a significant assumption: that each entity has been assigned exactly one key, which is difficult to ensure in practice in general, as we discuss below.

Our methods of detecting Sybil attackers are related to malicious attacks against anonymous routing protocols [19] called *intersection attacks* [1, 26, 31]. Anonymous routing protocols allow an identity to remain indistinguishable from other nodes in the system. An attacker that wishes to determine the identity of an initiator can track the membership of the group over time. Each time the attacker identifies a message, it records the group membership. As membership changes due to nodes joining or leaving the group purposely or because of network failures, the intersection of all the recorded memberships converges to only the initiator. Our work in this paper is an application of the intersection attack applied to geographic location in an ad hoc network.

Similarly, a Sybil attacker wishes to keep her multiple identities are indistinguishable from others in the system. However, there are differences between a Sybil attacker and legitimate nodes in a mobile wireless scenario, particularly in that independent nodes are mobile but the identities of a Sybil node move together. This provides the opportunity to identify a Sybil attacker using a location-based intersection mechanism because the Sybil identities will always be part of the intersected group.

In the remainder of this section, we present an overview of the security problems that the Sybil attacks pose to ad hoc networks in particular.

2.1 Sybil Attacks in Ad hoc Networks

An ad hoc network is composed of mobile, wireless devices, referred to as *nodes*, that communicate only over a shared broadcast channel. An advantage of such a network is that no fixed infrastructure is required: a network for routing data can be formed from whatever nodes are available. Nodes forward messages for each other to provide connectivity to nodes outside direct broadcast range.

Ad hoc routing protocols are used to find a path end-to-end through the cooperative network [25, 14]. Each node needs a unique address to participate in the routing. Often addresses are assigned as an IP addresses or a unique *media access channel* (MAC) address. Because all communications are conducted over the broadcast channel, nothing but these identifiers are available to determine what nodes are present in the network.

In unsecured routing protocols, such as DSR or AODV, these address-based identifiers can be easily falsified by malicious nodes, which presents an opportunity for a Sybil attack. However, allowing unauthenticated address presents a series of other attacks, including route direction, spoofing, and error fabrication [12]. Our methods work whether addresses are authenticated or not, though given the wide range of attacks possible against unauthenticated networks, Sybil attacks may not be the most significant problem present. Our methods will also work on disruption tolerant networks (e.g., [6]), however, just as such networks incur an extreme routing delay, there will be a corresponding large delay in successful sybil attack detection.

Secured ad hoc networks can be classified into three broad groups, each of which can be susceptible to the Sybil attack.

- **PKI-based protocols.** Much of the initial work in ad hoc network security focuses on secure routing [12, 28, 13, 11, 24, 23]. A variety of protocols have been proposed to counter routing attacks, some of which require a central authority or other mechanism to distribute cryptographic material to nodes in the system prior to or during deployment. Systems involving a central authority are less flexible, and installing a central authority removes the chief advantage of ad hoc networks: the ability to form spontaneously from whatever nodes are available. Allowing nodes to join without pre-distributing keys leaves a potential Sybil attack.

- **Threshold-based protocols.** To avoid the untenable requirement of a PKI, other protocols use *threshold cryptography*. In such scheme, a group of trusted nodes distributes cryptographic material only if a subset of that group agrees on the trustworthiness of new members [16, 32, 15]. Sybil attackers can additionally defeat schemes that rely on threshold cryptography because verifying the true number and independence of nodes in the network is difficult. If a Sybil attacker can generate identities to meet the threshold requirements it can effectively control the routing of the network.

- **Reputation Schemes.** Other security mechanisms for ad hoc networks include protocols for determining and maintaining reputation information about nodes in the group [3, 18, 2, 21, 27]. Each node can develop trust in the other nodes that it believes are routing correctly. The Sybil attack undermines these protocols because a node can use multiple identities to falsely vouch for or otherwise support an identity that would otherwise gain a bad reputation.

A reliance on cryptographic certificates or keys does not prevent the Sybil attack in general because one entity may be in possession of multiple keys. For example, if PKI credentials are simply purchased (e.g., through VeriSign), the PKI is reduced to a resource test of each identity's wealth, which can be without bound. Unfortunately, implementing a stronger approach is problematic. This is because in practice it is untenable to create a foolproof system that can scale to a significant number of users to check identities for independence before the keys are issued. Deploying a foolproof systems touches on issues including physical security and attacks involving social engineering or physical force. It would require checking a person against some set of unforgeable documents; but even government-issued documents are forged regularly.

3 Detecting the Sybil Attack

The identities established by a Sybil attacker — whether represented by IP addresses, MAC addresses, or public keys — differ from those of an honest node in several ways. Because the resources of a single node are used to simulate multiple identities, any particular assumed identity is resource constrained in computation, storage, or bandwidth. Douecer has shown that a Sybil attacker cannot be *prevented* by tests of finite resources [10]. However, unlike separate entities, all identities of a Sybil attacker must share the same set of resources, and this sharing can be *detected* in some scenarios [22].

In the mobile environment, a single entity impersonating multiple identities has an important constraint that can be detected: because all identities are part of the same physical device, they must move in unison, while independent nodes are free to move at will. As nodes move geographically, all the Sybil identities will appear or disappear simultaneously as the attacker moves in and out of range. Assuming an attacker uses a single-channel radio, multiple Sybil identities must transmit serially, whereas multiple independent nodes can transmit in parallel. The latter two differences form the basis of the Sybil attack detection scheme proposed here.

A single attacker with multiple radios will not be identified as such; however, there may be analogous methods for detecting multi-radio Sybil attackers, which we discuss in Section 6.

3.1 Overview

In our scheme, individual nodes that wish to detect Sybil attackers monitor all transmissions they receive over many time intervals. These intervals are chosen long enough to capture behavior from all the Sybil identities of an attacker, including data transmissions, HELLO and keep-alive messages, and routing requests and replies. The node keeps track of the different identities heard during the interval. Having made many observations, the node analyzes the data to find identities that appear together often and that appear apart rarely. These identities likely comprise a Sybil attack.

Like all detection systems this protocol produces accurate and desirable *true positive* results when it determines that a particular identity is part of a Sybil attack and *true negative* results when it correctly labels a independent node as such. It produces *false positives* by identifying an independent node as being part of the Sybil attack and *false negatives* by concluding a Sybil identity is just an independent node.

Each false result has several potential causes. False negatives can be caused by insufficient sampling time, so that transmissions from all the Sybil identities are not captured in each interval. Similarly, if the Sybil attacker does not transmit using all its different identities within an interval the results will be skewed. Sybil attackers may actively thwart detection by changing identities frequently. However, doing so limits the effectiveness of an attack when false identities would best be long-lived, for example to foil a reputation scheme or to defeat threshold cryptography.

False positives can be caused by using Sybil identities that belong to other nodes in the network. An attacker can corrupt trust in legitimate nodes in this way. False positives can also occur if a collection of nodes moves together in unison in close proximity, either accidentally or intentionally. For example, a military unit with many members maneuvering in formation, each of whom has a wireless device, will appear as a Sybil attacker based on their physical proximity. We show in Section 5 that we can reduce the rate of false positives in these cases by analyzing the rate of packet collision at the MAC layer.

Assumptions Several assumptions underly our protocol design. The first is that it should run on any normal node and not require any unusual hardware, nor does it require any directional antennas or specialized clocks. The protocol requires only that a node be able to receive transmissions. Later, when the protocol is extended to multiple nodes, it also requires they be able to share data among them by forwarding it, and that each node have a similar notion of what time it is to within a few seconds, both of which are part of the normal operation of the network. Because the hardware simplicity is low, the protocol can easily be widely deployed. Secondly, we assume that the Sybil attacker maintains its identities over time rather than disposing of them and creating new ones; that is, it is a simultaneous Sybil attacker [22]. This assumption is reasonable for Sybil attackers wishing to thwart long-lived protocols, such as those that use threshold cryptography or maintain reputation information. Finally, for the purposes of this paper, we confine discussion to networks with only a single Sybil node. We are confident this protocol can adapt to discover multiple attackers; however, we leave this as future work.

3.2 Detection Protocol

In this section, we describe two versions of our first detection protocol: a single observer case and a multi-observer case. In the next section, we evaluate both of these protocols. In Section 5, we show how these basic methods can be extended to include information from the MAC network layer.

3.2.1 Single Node Observer

Our protocol, Passive Ad hoc Sybil Identity Detection (PASID), is purely passive; it does not require active probing of suspected Sybil nodes, though the techniques are complimentary to active methods [22]. Instead, it operates effectively on a single node that records the identities of nodes that it hears broadcasting. In Section 3.2.2, we show how to improve protocol performance by sharing this data among a set of trusted nodes.

A node that wishes to detect the presence of a Sybil attacker in the network starts by recording the identities of all other nodes it hears broadcasting over a series of intervals. A complete record of all data transfered is not needed: it is sufficient to record only identities of the nodes heard during each interval. The observation period, referred to as the *time bucket*, is long enough to capture the likely behavior of a normal node including: normal data flow; regular HELLO and keep-alive messages; and periodic route requests for nodes that have data to send but have no current route to the destination. The length of this time period depends on the underlying protocol employed; in our simulations 30 seconds was adequate as it far exceeded the period between routing updates and requests. A more thorough investigation of bucket times would reveal the advantages of longer or dynamically chosen bucket times; however, we reserve this topic for future work.

After a sufficient observation period, which consists of a number of buckets, the node attempts to determine if it has observed a Sybil node. The length of the observation period depends on the amount of mobility within the network; highly mobile networks need fewer intervals than more static networks. In our simulations, 200 observations over 6,000 seconds, or 100 minutes of simulated time, was sufficient.

The node then determines which pairs of nodes are related. While correlation would be the most obvious candidate for doing so, it suffers from the fact that nodes that are never seen together will be highly correlated. In this case, however, nodes that are not seen together cannot be assumed to be related; they might be in separate parts of the network, unheard by both the observing node and each other. We therefore tried a number of different techniques to measure the relationship between nodes, including machine learning tools [30].

Our final and simple solution reflects the intuition that, during some observation period, seeing a pair together provides some evidence they are related; that seeing one but not the other of a pair provides stronger evidence they are not related; and that not seeing either of a pair nodes provides no evidence, because it is not possible to tell if they appear together elsewhere or separate elsewhere. Our solution also reflects that having more observations of the nodes in question provides more evidence than fewer observations.

After a period of observation, the detection algorithm then works in a series of simple steps:

1. We calculate A_{ij}, the affinity between nodes i and j, as
$$A_{ij} = (T_{ij} - 2L_{ij}) \frac{T_{ij} + L_{ij}}{N} \quad (1)$$
where T_{ij} is the number of intervals in which nodes i and j were observed together, L_{ij} is the number of intervals in which either i or j were observed alone, and N is total number of intervals in the observation period.

2. After the affinity between each pair of nodes has been computed, the observer constructs a graph in which the node identities are the vertices and the undirected edges are weighted with the affinity values between them. Only edges that are greater than a specific *threshold* parameter are included. Using our measure of affinity, we recorded our results using a threshold of 0.1.

3. Depth-first search (DFS) is then run over each vertex to discover the connected components. Each of the components found represents a possible Sybil attacker. While there can be several different connected components, we took only the largest to be a Sybil attacker, in

line with the working assumption that there was only one per network. If there were more, they would appear as separate components.

Note that this approach is clearly not optimal; DFS can have a long running time for large numbers of nodes. We will look to improve the scalability of the analysis algorithm in our future work.

The justification for this affinity measure is that each identity of an attacking node must transmit often enough to participate in the protocols that operate the network, including routing. If the observation periods are long enough, the attacker will be forced to transmit within a single period to maintain the fiction of multiple identities. For example, in AODV, routes that are not used for 3 seconds are dropped, thus our observation period is set to 30 seconds in our evaluations to catch route re-formations. Accordingly, we expect that for most realistic scenarios, the attacker will find it difficult have identities transmit individually in separate periods. In situations where this is not the case, the period can be lengthened or the weights of observations together and apart can be adjusted to account for the change in difficulty.

3.2.2 Multiple Node Observers

Though observations from a single node can accurately identify a Sybil attacker, any single observer is inherently limited in the area that can be monitored. Collaborating observers are not only able to cover a larger area, but might be able to determine that different identities are not related because they are seen in different areas at different times. PASID should, therefore, increase in accuracy as we add observers to the network.

We assume a subset of the legitimate nodes in the network can share observations periodically using the normal data transmission capabilities of the ad hoc network, and that these nodes can trust each other to perform this task honestly. Each node again tracks all other nodes that it hears over many time buckets. At the end of the observation period, it exchanges the information of what identities were heard during what time periods with the other nodes it trusts in the calculations. Note that this exchange does not have to occur often; in our simulations, it would only happen every 100 minutes. We do not simulate the exchange in our simulations, and assume that there is sufficient connectivity for all trusted nodes to reach each other; if this is not the case, detection will be delayed until node movement allows one or more nodes to accumulate the results of all observations. We will see, however, that the accuracy increases even if only some of the additional observations are received.

When using observations from more than one node, the counts are totals over all observing nodes and the last term of Equation 1 is at most 1. We let G represent the number of nodes sharing observations with one another. We modify our other variables to account for the multi-observer case. Now, $T_{ij}(n)$ is number of intervals in which nodes i and j were observed together by node n, defined as

$$T_{ij} = \sum_{n \in G} T_{ij}(n).$$

We let $L_{ij}(n)$ is the number of intervals in which either i or j were observed alone by node n, defined as

$$L_{ij} = \sum_{n \in G} L_{ij}(n).$$

N is still the total number of intervals in the observation period. Accordingly, the affinities for the multi-observer case are calculated as follows.

$$A_{ij} = (T_{ij} - 2L_{ij})\, w_{ij}$$

where

$$w_{ij} = \begin{cases} \dfrac{T_{ij} + L_{ij}}{N} & \text{if } T_{ij} + L_{ij} < N, \\ 1 & \text{otherwise} \end{cases}$$

4 Evaluation

To demonstrate the effectiveness of our first detection protocol, we simulate a series of ad hoc networks using the ns2 network simulator [20]. We evaluate the single-observer and multi-observer cases. Following this section, we introduce a modification to the protocol that makes use of MAC layer information. We show that a single node can accurately identify the various identities of a Sybil attacker, and that cooperating nodes can increase the accuracy of the process.

Kotz, et al [17] have expressed concerns about using ns2 to simulate mobile networking. However, we believe the evaluations we present below are sufficient as a preliminary exploration of our method. This is because our sybil detection approach relies on simple observations and is not dependent on the actual throughput or bit error rate of a channel or the particular efficacy of some routing or MAC protocol. In other words, at a minimum, our simulations clearly show the feasibility of our approach; more realistic models or evaluations over real traces would refine performance numbers. Additionally, we model mobility using the random way point model, which has also come under scrutiny. However, this model is reasonable here because it does not restrict mobility along some path or sub area for any particular node. Changing the mobility model would not affect operation of the protocol, though it would again refine the results.

4.1 Evaluation Assumptions

Each node in the network sends traffic to one other randomly selected node using constant bit rate traffic at 10Kbps. The Sybil node uses several 10Kbps connections — one for each node it tries to impersonate. Since the Sybil attacker is

Simulation Parameter	Value	Simulation Parameter	Value
Number of nodes	5, 10, 25, 40	ns parameter: netif	Phy/WirelessPhy
Topography size (m)	250, 500, 1000, 2000, 3000, 4000	ns parameter: mac	Mac/802.11
		ns parameter: ll	LL
Number of Sybil identities	5, 10, 20	ns parameter: ifqlen	50
Simulation time (sec)	6000	ns parameter: traffic type	level
Time bucket length (sec)	30	ns parameter: packet size	512
ns parameter: chan	WirelessChannel	ns parameter: packet interval	.25
ns parameter: prop	TwoRayGround	ns parameter: packet max	1,000,000
ns parameter: seed	0.0	ns parameter: start time	1.0
ns parameter: adhocRouting	AODV	ns parameter: pause time	10
ns parameter: ant	Antenna/OmniAntenna	ns parameter: max speed	5
ns parameter: ifq	DropTail/PriQueue		

Table 1: Simulation Parameters

simulated using data streams, routing and control packets are under-represented in our simulations. Therefore, our simulations represent worst-case results. Although we use the AODV [25] mobile ad hoc routing protocol, we ignore HELLO messages because not all ad hoc routing protocols use them; to count them would skew the results. In practice, a Sybil attacker would have to accurately simulate control traffic to avoid detection based on aberrant behavior, and PASID would be more effective than shown here.

Our results are averaged over many simulation runs. We create ad hoc networks in which nodes wander in square topographies. Topography size is given by the length of the square's side. The number of legitimate nodes varies between 5, 10, 25, and 40, and the number of identities assumed by the Sybil attacker, in those experiments with an attacker, varied between 5, 10, and 20. The simulations each run 6,000 seconds and the duration of each observational time bucket is 30 seconds. Each combination of the above parameters was repeated 20 times for each topography, for a total of 1,440 runs that took over a day to complete on a 32-node computing cluster. Other simulation parameters are show in Table 1.

4.2 Single Node Observer

Figure 1(a) shows the rate at which a single node observer will falsely identify other nodes as Sybil attackers when Sybil nodes are present, averaged over all experiments for each topography size. The true positive rate is zero because there are no Sybil identities present. For topographies smaller than 1000m-by-1000m, the false positive rate is high because all nodes are frequently within range of each other. Because of this, nodes will be very frequently heard together even when they are not Sybil identities, and will rarely be heard apart as they do not move out of radio range. This leads to the false identification rate in topographies that are denser in terms of nodes per square meter.

The accuracy and error rates for a single node observer when a Sybil attacker present is shown in Figure 1(b). Again, in smaller topographies there is insufficient mixing to separate Sybil identities from real nodes, and the error rate is high, as is the detection rate, because all nodes are seen as part of the same identity. As the topography size increases, the number of meaningful observations that a single node can make increases, and the true positive rate stays high, on the order of 95%, while the false positive rate drops significantly. As the topography size increases further, the number of observations that a single node can make is reduced as all nodes are spread far apart, and the accuracy of identifying the Sybil identities decreases.

It is important to note that this is a single node, acting alone, using only minimal hardware in a passive manner. This initial result indicates that the technique can work, and work well in some circumstances. As we show in the following sections, the accuracy increases as we add additional observers.

4.3 Multiple Observers

To test the effectiveness of multiple observers sharing data, we used the data from the simulation runs shown above, but we calculated the detection accuracy for different sizes of collaborating groups of nodes. Because the number of non-attacking nodes in the simulations vary between 4 and 39, we measured the number of collaborating nodes as a percentage of group membership instead of an absolute number. The graphs in Figure 2 show the false positive and true positive rates for increasing percentages of collaborating nodes. The error bars, representing the 95% confidence interval, are larger than the graphs in Section 4.2 because the results shown are averaged over fewer number of simulation runs: as it was computationally prohibitive to test all possible sets of collaborating nodes, we instead randomly chose 10 different groups for each measurement.

The results show that adding independent observers in the network can significantly increase the accuracy of identifying the identities being used by a Sybil attacker. With at least 20% of the nodes in the network sharing information, it is in some cases possible to achieve 100% accuracy with 0% false positives. This sweet spot lies between the area in which there is not enough mixing of nodes to do detection, and the area in which the network is sparse that nodes do not encounter each other frequently enough to make meaningful observations. Our simulation data shows that this is increase in accuracy is true regardless of the number of nodes in the network.

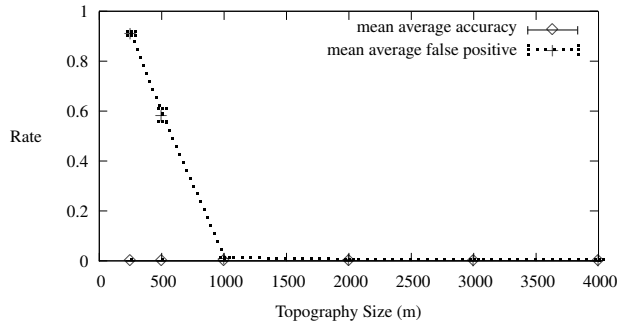

(a) No Sybil attacker present

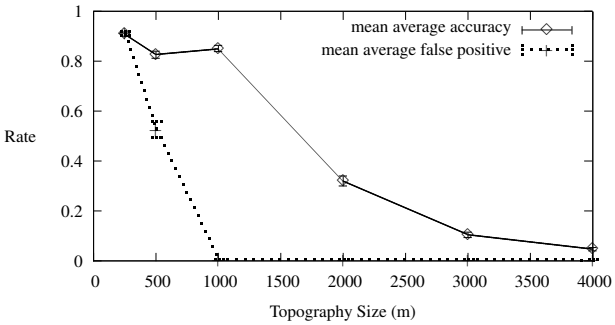

(b) Single Sybil attacker present

Figure 1: Accuracy rates vs. topography size for a single observer

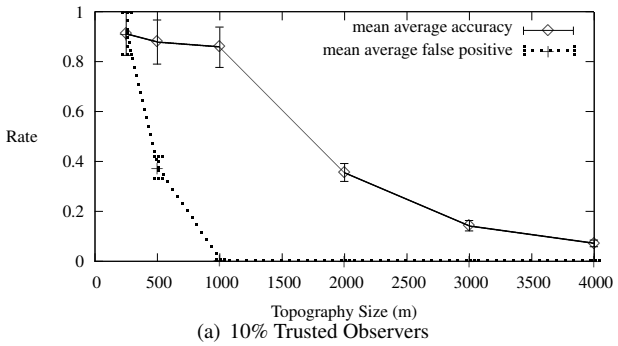

(a) 10% Trusted Observers

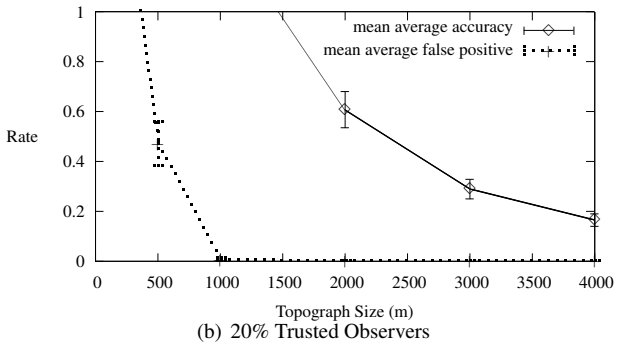

(b) 20% Trusted Observers

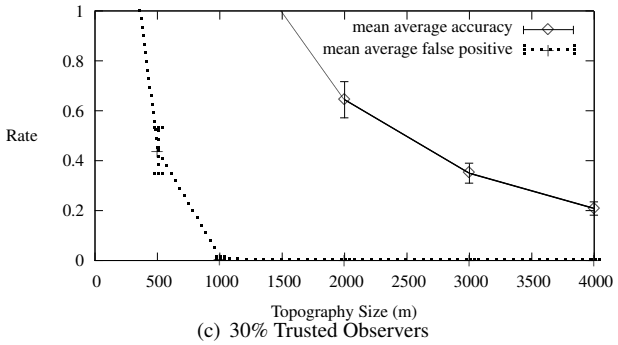

(c) 30% Trusted Observers

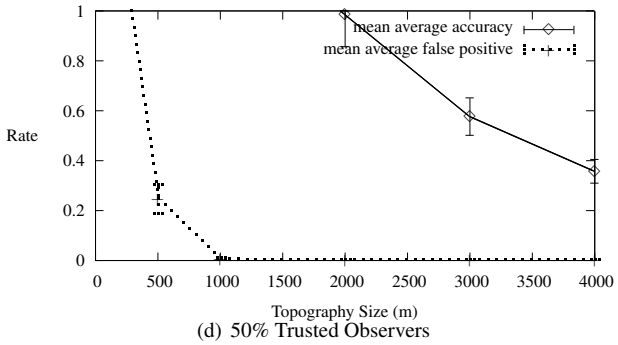

(d) 50% Trusted Observers

Figure 2: Accuracy rates for various percentages of observers

4.4 Discussion

In this initial approach, there are a number of parameters that must be determined in order for PASID to function effectively, including how long the sample period should be, how many are needed, and what the threshold should be. These parameters will vary depending on the conditions of the individual network. In this work, we determined these on an informal basis by trying different values until we had what seemed to be reasonable results. A better approach might be to have a trusted node openly perform the Sybil attack itself, and then to use the data collected from that experiment to tune the parameters to optimize detection. Combined with receiver operating characteristic (ROC) analysis to tune the parameters to the specific network, this could potentially increase detection rates.

The technique of combining data from many different trusted observer nodes also currently requires that all nodes act correctly; the formula used to determine the relationship

between is very prone to an increased error rate if one of the trusted nodes behaves maliciously. Our future work includes investigating other formulae that better reflect statistical models and that are more resistant to malicious input.

This work also assumes that the Sybil attacker maintains the same identities over the entire observation period. We believe this to be a reasonable assumption, particularly for reputation schemes. However, a Sybil attacker that suspects it is being monitored might try a variety of things to thwart this process. This might include not using some of its identities at different times; changing identities over time; or using the identities of existing nodes at different times, if possible. Our future work includes an examination of the effectiveness of these techniques and how PASID might be improved to deal with them.

5 Sybil attackers and proximate groups

PASID is predicated on the fact that all identities of a single Sybil attacker must appear to move together over time, while other independent nodes move individually. However, there are clearly cases in which a group of distinct mobile nodes might move together in unison. For example, a military unit that was equipped with mobile nodes might all move as a group for long periods of time. In cases like these, a group moving in close proximity will be detected as a Sybil attacker, and each individual node will be mistaken for a Sybil identity.

There is a detectable difference between nodes moving together and a Sybil attacker, however. A Sybil attacker with a single radio can only transmit messages for all its identities in serial, while a group of radios will be able to transmit in parallel. While it would be possible to detect this by active probing of the suspected Sybil attacker's resources, it is also possible to detect this passively. A group transmitting in parallel will cause collisions at the multiple access channel (MAC) level when different nodes try to transmit simultaneously. Serial transmissions from a Sybil attacker will result in fewer such collisions.

In this section, we first show that groups traveling together are accurately detected by PASID. We then show that groups generate higher MAC collisions rates than Sybil attackers, and that we can use the collision rate to differentiate between a group moving together and a Sybil attacker.

5.1 When nodes move together

Nodes that move in close proximity will be detected as a Sybil attacker using PASID. To show this, we ran the same simulations as in Section 4; however, we replaced the Sybil attacker with a group of nodes moving in close proximity. The number of nodes in the group was the same as the number of different identities being presented by the Sybil attacker. The nodes were spaced one meter apart, and placed randomly within the smallest square possible centered on where the Sybil attacker moved. For example, a group of 20 nodes was placed randomly within a 5 by 5 meter area, 1 meter apart. The group maintained their relative positions as the followed the same path taken by the Sybil attacker in the prior simulations.

Figure 3 shows the results of the simulations with groups in place of Sybil attackers. Comparison with Figures 1 and 2 demonstrates that groups are detected with approximately the same accuracy as a Sybil attacker. Given the goal of detecting a Sybil attacker, this means that any legitimate group will generate a false alarm. We discuss how to avoid this in the next section.

5.2 Differentiating between groups

A single Sybil node differs from a group of individual nodes in two ways. First, its identities must be located physically in the same place. Second, it can only transmit messages serially. We have shown in Section 4 that it is possible to identify all individual Sybil identities based on their location; in this section, we show that it is possible to differentiate between a Sybil attacker and a group moving together by detecting whether transmissions occur serially or not.

While it is possible to test for serial transmissions actively [22], we follow our low-overhead and passive approach, and instead monitor collisions at the MAC level. Because it is a shared broadcast channel, individual nodes do not coordinate when messages are sent. Instead, they each send data as it becomes available, and on occasion individual senders sometimes start transmitting simultaneously. This is called a collision, and when it occurs neither message is received because the two transmissions interfere with each other. The intuition behind our second method, named PASID-GD for PASID with Group Detection, is that a Sybil attacker, limited to serial transmissions, will cause fewer collisions than a group of independent nodes transmitting in parallel.

To test this, we again used the simulation data from Section 4. To verify the intuition, we first measured the number of MAC collisions for each case: *base*, which is our control group with no Sybil attacker or group moving together; *Sybil*, which includes a single Sybil node; and *echelon*, which is a group moving together. Figure 4(a) shows the average number of collisions per time bucket for each of these cases. It is clear that there is a significant difference between the Sybil and echelon cases, with far more collisions occurring when a real group is present.

Knowing that there is a difference between the two cases, we can develop a method to detect it. However, because an individual node can only operate in the basis of the transmissions it witnesses, no node can know what the base rate of collisions would be if the group was not present. Instead, PASID-GD of detection operates by measuring and compar-

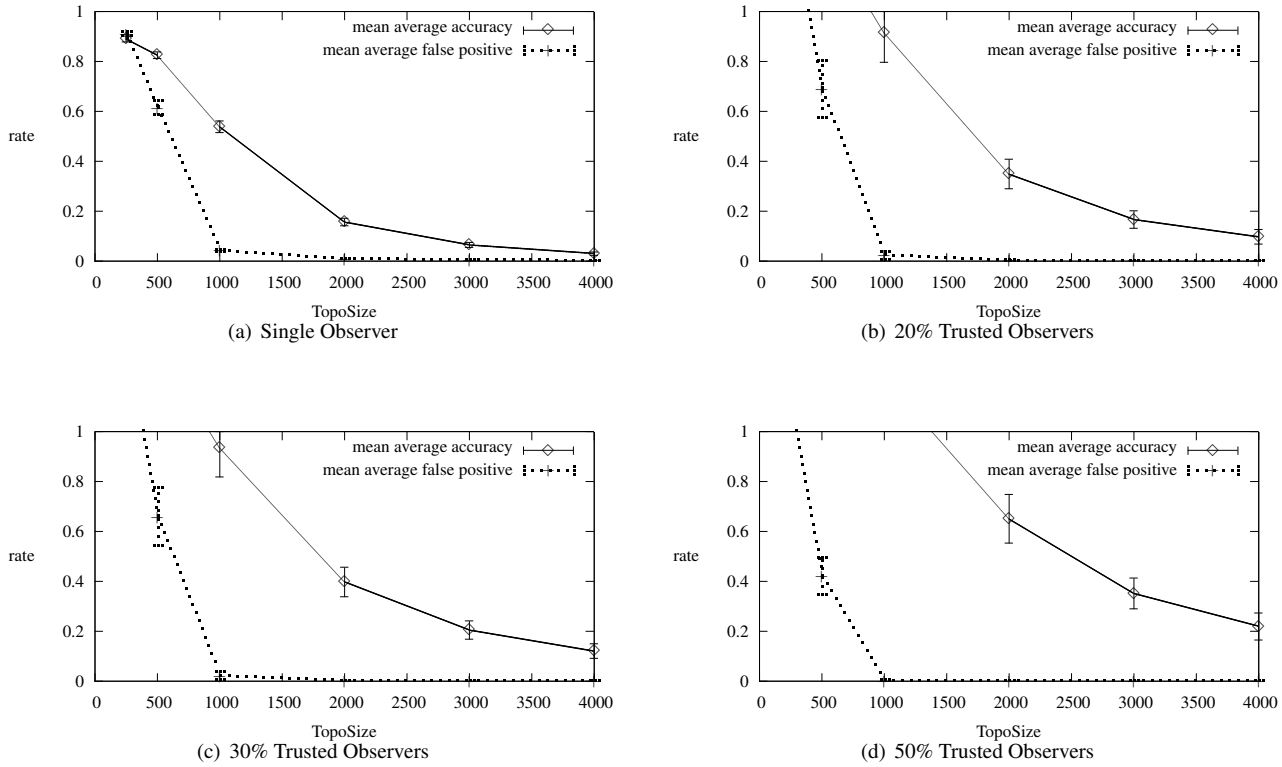

(a) Single Observer (b) 20% Trusted Observers

(c) 30% Trusted Observers (d) 50% Trusted Observers

Figure 3: Accuracy rates for various percentages of observers, Group moving in Unison

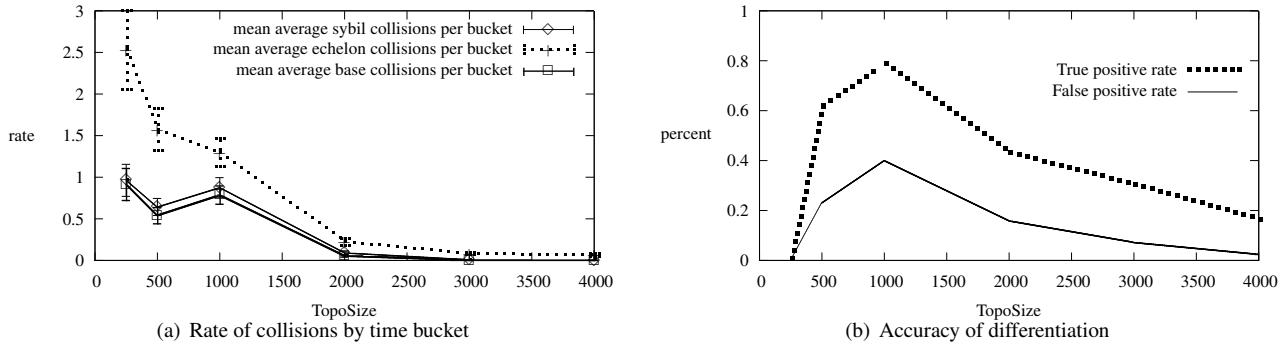

(a) Rate of collisions by time bucket (b) Accuracy of differentiation

Figure 4: The use of MAC collisions to differentiate Sybil attackers from groups

ing the rate of collisions when a detected group is present to when it is absent. This requires keeping a count of all collisions received during each time bucket.

After a group of identities is detected based using PASID, each node determines an average collision rate for the time periods when the group was present and when the group was not present. It then compares the ratio of the two. If this is above a particular threshold, 1.5 in our simulations, then it assumes that it is a group moving together; otherwise the node assumes it is a Sybil group.

Figure 4(b) shows the average accuracy and false positive rates of this approach. The average accuracy is the percentage of times that a Sybil group is correctly identified as such, and is close to 80% in some conditions. The false positive rate is the percentage of time that a group moving together is identified as a Sybil node, and is always less than 40%. Topography size has similar effects here as it does on PASID. When the topography size is small the density of nodes increases, and the base rate of collisions high enough that there is not a detectable difference when the

even a group moving together comes near. When the group becomes sparse, there are so few collisions on average that even a Sybil node causes enough to pass the threshold and appear as a separate group, despite the fact that there are fewer additional collisions than a group would cause. We are examining other methods of using MAC collisions to differentiate the two cases and are confident that the accuracy will increase.

6 Limitations and Future Work

We have shown in this initial work how individual mobile nodes with very common, inexpensive equipment can detect a Sybil attacker. This protocol is exceptionally simple and costs very little; a single node needs only a small amount of memory to record its observations. We also showed that trusted, collaborating nodes can greatly improve the speed and accuracy of identification by sharing their observations with each other periodically. Our longer term work is to investigate improvements in detection that more expensive computer hardware allows.

6.1 Extensions

PASID and PASID-GD are predicated on very simple hardware, and no special action other than recording and sharing observations on the part of participating nodes. Given the rapid advances being made in computer hardware, our longer term work will determine what other capabilities that could be common for both the attacker an observer.

The simplest boost to the protocol would be to have nodes move in a particular pattern to best cover the entire network and to create more opportunities for nodes to move in and out of range. Our experiments were carried out with nodes moving in a random pattern, which allowed some mixing, but not as much as deliberate motion would. Multiple cooperating nodes could organize their motion to better cover the entire area.

If mobile nodes cooperating in detecting a Sybil node had more accurate and closely synchronized clocks than what we have assumed, then it is possible for the trusted nodes to establish their relative positions using one of a large number of localization schemes [5, 7, 4, 29]. Once the relative position was determined, then each node could record the time that every individual message was received. By then comparing these times , each node could determine the distance from each sender. This could quickly show that different nodes were not close together, ruling out the possibility that they were a Sybil attacker.

As hardware costs decrease, it is likely that many mobile nodes will contain Global Positioning System (GPS) receivers. This will let nodes know their position down to within a meter. Given just this absolute position information, nodes could collaborate to determine the position of an ad hoc Sybil attacker. In this case, each node could record its position when it received a message. Later comparison of the locations in which the same messages were received by different nodes would quickly indicate if different identities were heard in different locations that were far enough apart to rule out movement between them, indicating that they were not a Sybil attacker.

Finally, as mobile node receive additional antennas that allow from some measurement of signal direction, detecting Sybil attackers should become much faster. A mobile node can now record not only when it heard any identity, but the relative direction of the signal. All Sybil identities will show a correlation not only with position, but with relative distance, greatly reducing the false positive rate in many cases. Additionally, collaborating nodes could perform triangulation measurements to determine the location of many different transmitting identities; again, this would speed the discovery a Sybil attacker.

All the methods are completely passive in that they do not require any active probing of the suspected Sybil attacker. Instead, they can all be performed using passive observation and normal data communication to transmit their observations between observing nodes.

7 Conclusion

In this paper, we have presented the first completely passive approach to detecting a Sybil attacker in a network. PASID detects which network identities are related and likely to belong to the same Sybil attacker by monitoring what identities seem to be physically located together, and it can achieve 90% or more accuracy with no false positives in some circumstances. Adding additional observer nodes increases the accuracy to 100% and increases the range over which this accuracy is possible. We also have shown that PASID will detect a group of nodes moving together as a Sybil attacker, and we presented an extension to the method called PASID-GD that monitors collisions at the MAC level to differentiate between the single Sybil attacker and a group moving together.

We do not investigate how multiple collaborating Sybil attackers can thwart observers. We plan to address this in future work, and to investigate how different mobility models of attackers and observers can affect our results and the selection of our parameters, including using real-world mobility data. For the UMass DieselNet network [6], we deployed 802.11-based networking on 40 buses that roam a 150 square mile area 14 hours a day, and can use this data to model more realistic movement. We can also easily deploy this observation method on the buses for further experimentation.

Acknowledgments We would like to thank Prof. David Caraballo and Prof. Paul Kainen for helping experiment with various methods of determining node affinity.

References

[1] O. Berthold, H. Federrath, and M. Kohntopp. Project anonymity and unobservability in the internet. In *Computers Freedom and Privacy Conference 2000 (CFP)*, April 2000.

[2] S. Buchegger and J. Le Boudec. Performance Analysis of the CONFIDANT Protocol. In *Proc. Intl Symp on Mobile Ad hoc Networking and Computing*, pages 226–236, June 2002.

[3] S. Buchegger and J. Le Boudec. A Robust Reputation System for P2P and Mobile Ad hoc Networks. In *Proc. Wkshp Economics of Peer-to-Peer Systems*, June 2004.

[4] N. Bulusu, D. Estrin, L. Girod, and J. Heidemann. Scalable Coordination for wireless sensor networks: Self-Configuring Localization Systems. In *Proc. Intl Symp on Communication Theory and Applications*, July 2001.

[5] N. Bulusu, J. Heidemann, and D. Estrin. GPS-less Low Cost Outdoor Localization For Very Small Devices . *IEEE Personal Communications, Special Issue on Smart Spaces and Environments*, 7(5), October 2000.

[6] J. Burgess, B. Gallagher, D. Jensen, and B.N. Levine. Maxprop: Routing for vehicle-based disruption-tolerant networks. In *Proc. IEEE INFOCOM*, April 2006.

[7] S. Capkun, M. Hamdi, and J. Hubaux. GPS-Free Positioning in Mobile Ad hoc Networks. In *Proc. Hawaii Intl Conference on System Sciences*, 2001.

[8] S. Capkun, J. Hubaux, and L. Butty. Mobility helps security in ad hoc networks. In *Proc. ACM Intl Symp on Mobile Ad hoc Networking and Computing*, pages 46–56, June 2003.

[9] A. Cheng and E. Friedman. Sybilproof Reputation Mechanisms . In *ACM Wkshp on the Economics of Peer-to-Peer Systems*, August 2005.

[10] J. R. Douceur. The Sybil Attack. In *Intl Wkshp on Peer-to-Peer Systems*, March 2002.

[11] Y. Hu, D. Johnson, and A. Perrig. SEAD: Secure Efficient Distance Vector Routing for Mobile Wireless Ad hoc Networks. In *Proc. Wkshp on Mobile Computing Systems and Applications*, Jun. 2002.

[12] Y. Hu and A. Perrig. A Survey of Secure Wireless Ad hoc Routing. *IEEE Security & Privacy*, 2(3):28–39, May/June 2004.

[13] Y. Hu, A. Perrig, and D. Johnson. Ariadne: A Secure On-Demand Routing Protocol for Ad hoc Networks. In *Proc. Intl Conference on Mobile Computing and Networking*, Sep. 2002.

[14] D. Johnson and D. Maltz. Dynamic Source Routing in Ad hoc Wireless Networks. In *Mobile Computing*, volume 353. Kluwer Academic Publishers, 1996.

[15] A. Khalili, J. Katz, and W. A. Arbaugh. Toward Secure Key Distribution in Truly Ad hoc Networks. In *Proc. Symp on Applications and the Internet Wkshps*, January 2003.

[16] J. Kong, P. Zerfos, H. Luo, S. Lu, and L. Zhang. Providing Robust and Ubiquitous Security Support for Wireless Mobile Networks. In *Proc. Intl Conference on Network Protocols*, Nov. 2001.

[17] D. Kotz, C. Newport, R. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *Proc. ACM/IEEE Intl Symp on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, pages 78–82, October 2004.

[18] S. Marti, T. J. Giuli, K. Lai, and M. Baker. Mitigating Routing Misbehavior in Mobile Ad hoc Networks. In *Mobile Computing and Networking*, pages 255–265, 2000.

[19] N. Mathewson, P. Syverson, and R. Dingledine. TOR: The Second-Generation Onion Router. In *Proc. USENIX Security Symp*, August 2004.

[20] S. McCanne and S. Floyd. Network Simulator Version 2. `http://www.isi.edu/nsnam/ns`.

[21] P. Michiardi and R. Molva. Core: A Collaborative Reputation Mechanism to Enforce Node Cooperation in Mobile Ad hoc Networks. In *Proc. Communications and Multimedia Security*, September 2002.

[22] J. Newsome, E. Shi, D. Song, and A. Perrig. The Sybil Attack in Sensor Networks: Analysis & Defenses. In *Proc. Intl Symp on Information Processing in Sensor Networks*, 2004.

[23] P. Papadimitratos and Z. Haas. Secure Routing for Mobile Ad hoc Networks. In *Proc. Communication Networks and Distributed Systems Modeling and Simulation Conference*, Jan. 2002.

[24] P. Papadimitratos and Z. Haas. Secure Link State Routing for Mobile Ad hoc Networks. In *Proc. Symp on Applications and the Internet Wkshps*, January 2003.

[25] C. E. Perkins and E. M. Royer. Ad hoc On-Demand Distance Vector Routing. In *Proc. WMCSA*, Feb. 1999.

[26] J.-F. Raymond. Traffic analysis: Protocols, attacks, design issues and open problems. In *Proc. Intl Wkshp on Design Issues in Anonymity and Unobservability*, volume 2009 of *LNCS*, pages 10–29. Springer-Verlag, 2001.

[27] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman. Reputation Systems. *CACM*, 43(12):45–48, 2000.

[28] K. Sanzgiri, B. Dahill, D. LaFlamme, B. N. Levine, C. Shields, and E. Belding-Royer. A Secure Routing Protocol for Ad hoc Networks. *JSAC Special Issue on Ad hoc Networks*, March 2005.

[29] A. Savvides, C. Han, and M. Strivastava. Dynamic fine-grained localization in Ad hoc networks of sensors. In *Proc. international conference on Mobile computing and networking*, pages 166–179, 2001.

[30] I. Witten and E. Frank. *Data Mining : Practical Machine Learning Tools and Techniques with Java Implementations*. Morgan Kaufmann, 1999.

[31] M. Wright, M. Adler, B. N. Levine, and C. Shields. The Predecessor Attack: An Analysis of a Threat to Anonymous Communications Systems. *ACM TISSEC*, 7(4), November 2004.

[32] L. Zhou and Z. J. Haas. Securing Ad hoc Networks. *IEEE Network*, 13(6):24–30, 1999.