Countering False Accusations and Collusion in the Detection of In-Band Wormholes

Daniel Sterne, Geoffrey Lawler
SPARTA Inc.
{dan.sterne, geoff.lawler}@sparta.com

Richard Gopaul, Brian Rivera, Kelvin Marcus
U.S. Army Research Laboratory
{rgopaul, brivera, kmarcus}@arl.army.mil

Peter Kruus
The Johns Hopkins University Applied Physics Laboratory
peter.kruus@jhuapl.edu

Abstract

Cooperative intrusion detection techniques for MANETs utilize ordinary computing hosts as network intrusion sensors. If compromised, these hosts may inject bogus data into the intrusion detection system to hide their activities or falsely accuse well-behaved nodes. Approaches to Byzantine fault tolerance involving voting are potentially applicable, but must address the fact that only nodes in particular topological locations at particular times are qualified to vote on whether an attack occurred.

We examine these issues in the context of a prototype distributed detector for self-contained, in-band wormholes in OLSR networks. We propose an opportunistic voting algorithm and present test results from a 48-node testbed in which colluding attackers generate corroborating false accusations against pairs of innocent nodes. The results indicate that opportunistic voting can instantaneously suppress false accusations when the network topology and routes chosen by OLSR provide a sufficient number of nearby honest observers to outvote the attackers.

1. Introduction

Detecting certain kinds of attacks on mobile ad hoc networks (MANETS), especially misbehavior in regard to routing protocols and packet forwarding, requires cooperative distributed intrusion detection techniques [17][23][24]. Such techniques utilize ordinary computing hosts as network intrusion sensors. Because ordinary hosts run complex applications, they are far more vulnerable to cyber penetration and compromise than the dedicated network intrusion detection systems typically used to monitor wired networks. If a host is compromised, it can inject bogus data into the cooperative intrusion detection system. This data can be used to conceal malicious activities, i.e., evade detection. Of greater concern is the potential injection of false accusations. These could stimulate an intrusion response system to curtail the access of falsely-accused victims, leading to a self-inflicted denial of service.

This form of the Byzantine fault tolerance problem is a critical, fundamental issue for any cooperative intrusion detection system. Well-known approaches to Byzantine fault tolerance typically involve some form of voting in which correctly executing replica processes are able to outvote a smaller number of dishonest (faulty) processes, often less than one third of the total population [15]. Applying such approaches to cooperative intrusion detection in MANETs confronts several obstacles. Because traffic in a MANET is dispersed throughout the network, each node has a unique view of network activity, so distributed intrusion sensors are not replicas. Furthermore, only nodes that are in particular topological locations can observe the attack traffic or network symptoms in question and reliably attest to whether or not an attack occurred. In other words, the set of nodes that is eligible to vote depends on the topology of the network and the routes used or affected by the attack. Moreover, the topology of a MANET may change continually; hence the set of nodes that is eligible to vote about a particular attack may vary over time.

We examine these issues in the context of an in-band wormhole attack on the OLSR routing protocol [14]. In this attack, a collection of colluding attackers creates the illusion that a single-hop "short cut", (a layer 2 link) exists between two nodes in distant parts of the network. The attackers create this illusion by connecting the purported neighbors by a much longer, covert tunnel through other unsuspecting nodes. If

Prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

well-positioned, the shortcut can attract a significant amount of traffic, diverting it from optimal routes so that it flows instead through the attacking nodes. This enables the attacker to delay, discard, or corrupt the attracted traffic at an opportune time. Another wormhole impact is that the attracted traffic travels along a highly suboptimal path. As a result, in-band wormholes consume network capacity (waste bandwidth). This inherent degradation of network service may additionally be accompanied by increased congestion, packet loss, and delay. In contrast, out-of-band wormholes, which have received far more attention, connect purported neighbors via a wireline network or additional RF channel that actually *adds* capacity to the network. For this reason, out-of band wormholes are arguably less harmful than in-band wormholes. Furthermore, because they do not require additional communications hardware, in-band wormholes may be more likely to occur in practice.

Building on our prior research [14], we describe a prototype intrusion detection system that detects and localizes wormholes by measuring and analyzing roundtrip delay times on multi-hop paths. Before concluding that an attack is underway, the detector requires agreement, by a majority of independent pairs of nodes, that a common link in these paths exhibits anomalous delays. This majority rule instantaneously suppresses false accusations even if they are corroborated by colluding attackers, as long as a sufficient number of honest nodes are appropriately positioned topologically. We refer to this approach as *opportunistic voting* because it utilizes as voters whichever nodes are in the right place at the right time. The use of voting to make end-to-end path measurements resilient to injection of bogus data appears to be unique to our detector. Nevertheless this notion may have broader potential value for increasing the Byzantine resistance of network path measurement techniques used in other intrusion detection approaches [10][12], as well as network fault localization [20] and network tomography [8].

Our paper makes the following contributions:

- We present an opportunistic voting algorithm for detecting and localizing in-band wormholes in which votes are based on roundtrip delay measurements reported by independent pairs of nodes. Voting is used for false accusation resistance.
- We describe a form of wormhole attack that includes coordinated false accusations by colluding pairs of nodes and actions by them to trick benign nodes into accusing other nodes.
- We present specialized tests and results showing the algorithm's effectiveness against a real implementation of this attack running in a network emulation testbed. The results show that opportunistic voting provides substantial value by instantaneously suppressing false accusations in dynamic topologies and at varying network densities.
- We discuss the effects of topology and routing on the eligibility of nodes to vote on the innocence of an accused link.

This paper is organized as follows. Section 2 discusses related work. Section 3 provides background on the in-band wormhole attack. Section 4 presents the wormhole detection and Byzantine resilience strategies we propose and the prototype detector we have constructed. Sections 5 and 6 describe prototype testing and test results; these are discussed further in Section 7. Section 8, the conclusion, provides a summary and identifies future directions.

2. Related work

A number of other researchers have developed intrusion detection techniques for wormhole attacks on MANETs. Their techniques generally differ from ours in two ways: they do not address false accusations, and they are concerned with out-of-band rather than in-band wormholes.

The concept of an out-of-band wormhole in ad hoc networks was introduced by Hu [9], who outlines temporal and geographic countermeasures designed to detect the remote forwarding of packets. Hu describes packet leashes, which attempt to restrict the maximum transmission distance of a packet. In this scheme, wormhole paths will cause neighbor-sensing and other packets to be received outside a tightly synchronized time window, so that the packets are treated as invalid.

Lazos describes a different geographic-based distance bounding approach for defeating out-of-band wormhole attacks [16]. With this approach, essential nodes are aware of their geographic location relative to specialized guard nodes. Messages forwarded beyond their local reach as determined by the guards, or detected as duplicates, are considered to be indicative of a wormhole. This approach may not provide adequate protection if the attackers are trusted insiders with access to the same node resources as other network nodes. Also, a highly mobile network may violate some of the assumptions made about node placement and density.

Adjih describes several approaches for countering out-of-band wormhole attacks on OLSR networks [3]. One technique is a geographic-based distance bounding approach for containing the relaying of HELLO and TC messages. Another is a watchdog-like

guard similar to [16] and [13] that attempts to monitor the number of packets sent and received by a given node.

Buttyán proposes techniques for detecting out-of-band wormholes based on statistical changes to neighbor hop counts and path lengths [7]. Local information is collected and reviewed by a central base station that detects changes in network behavior. These techniques are designed for static networks and do not identify the attackers.

Khalil describes MANET wormhole defensive countermeasures for routing protocols like DSR and AODV [13]. These are based on specialized guard nodes that promiscuously monitor all forwarding paths, an approach resembling the distance-bound approach proposed by [16].

The research we present here builds on our prior work, which provided the first detailed description of in-band wormholes in OLSR networks [14] [11]. We introduced the terms *extended wormhole* for attacks in which the wormhole link appears between unsuspecting benign nodes and *self-contained wormhole* for attacks in which the wormhole link appears between two attacking nodes. This work provided initial evidence that in-band wormholes can be detected by measuring roundtrip loss and delay. It also defined metrics to characterize the impact of wormholes on the surrounding network [11].

Awerbuch proposes the On-Demand Secure Byzantine Routine protocol (ODBSR), and describes its ability to defend against various attacks, including out-of-band wormholes [2]. Awerbuch's "Byzantine wormhole" is the out-of-band equivalent of our self-contained in-band wormhole. ODBSR is evaluated using simulations including a rectangular mesh of interconnected wormhole links called a "superwormhole". ODBSR mechanisms do not detect wormholes per se; they detect packet dropping that is applied to traffic traveling through wormholes. If a wormhole is simply attracting traffic while waiting for an opportune time to disrupt it, or is already subjecting that traffic to delays or packet corruption, the wormhole will remain invisible to ODBSR.

Gorlatova describes the detection of out-of-band wormholes in OLSR networks using an approach based on the same underlying premise as ours [10]: the path characteristics of a wormhole link should have a measurable effect on traffic flowing across it. Gorlatova, however, applies signal processing techniques to periodic incoming messages having a known frequency, in this case, OLSR HELLO packets. Each node analyzes the distribution and power spectral density (PSD) of inter-HELLO arrival intervals from each of its neighbors. If the HELLOs have arrived via a wormhole tunnel, the associated delay, even if quite small, is said to broaden or smear the HELLO message time series.

Ilsam proposes a system for suburban ad hoc networks that can detect violation of MAC and bandwidth reservation protocols, dropping packets, and delaying packets. [12]. Although Ilsam does not discuss wormholes, two of his techniques are related to ours. The first monitors hop-by-hop transit times for flows to identify nodes responsible for unusual delays. The second measures roundtrip delays using probe packets, which are assumed to be stealthy to intermediate nodes.

Intrusion detection techniques for other kinds of attacks on MANETS have given more attention to false accusations than those for wormholes. For example, Buchegger proposes a reputation system for MANETs in which nodes collect first-hand information about other nodes by direct observation and share it with selected other nodes, e.g., their neighbors [5]. Each node maintains a rating (a continuous variable) for every other node that it cares about. The system is designed to address two kinds of misbehavior: a) routing or forwarding misbehavior, and b) lies about the reputation of other nodes. Using second-hand information about both good and bad behavior of others is said to accelerate the detection and isolation of malicious nodes, but can endanger the integrity of the reputation system because nodes can use it to disseminate misinformation. This system improves upon Buchegger's earlier reputation system called CONFIDANT [4] which was susceptible to false accusations. To counter this problem, each node accepts second-hand information only if it is not too dissimilar from that node's first-hand information.

In Buchegger's and most other reputation systems [4][5][25], each node can maintain its own rating of other nodes because it bases that rating on *observations it can make independently*. For example, if one node forwards a packet to another node, it may be able to use promiscuous monitoring to determine whether the receiver continued forwarding the packet or intentionally dropped it[1]. By contrast, end-to-end path probing, which is useful for detecting in-band wormhole attacks and other kinds of routing misbehaviors, and for localizing network faults, minimally *requires that pairs of path endpoints cooperate to produce usable measurements*. Hence end-to-end path observations cannot be made by nodes acting independently.

[1] Marti describes a number of conditions under which promiscuous monitoring is unreliable [MGLB00].

3. In-band wormholes in OLSR

In OLSR, a proactive link-state routing protocol [21], the status of 1-hop links is gathered through the exchange of HELLO messages among 1-hop neighbors. Topology Control (TC) messages are then used to propagate link-state information to all other nodes. From this information, nodes formulate next-hop routing decisions based on the shortest-path computations.

In an in-band wormhole attack, two distant colluding nodes create a shared tunnel and use it to covertly transfer OLSR control messages (e.g., HELLO and TC messages) between remote nodes. The eventual result is that OLSR is tricked into thinking that the remote nodes are one-hop neighbors and that the link between them is a shortcut. This attracts traffic, and the attracted packets are then forwarded by the attackers through the tunnel, completing the wormhole illusion.

An in-band wormhole can fall victim to its own success, as the disruption in network routing can also affect the routing of traffic through the wormhole tunnel, causing the wormhole to collapse [14]. An in-band wormhole collapses when its tunnel endpoints cannot continue to forward control messages between remote network regions. One way to avoid tunnel collapse is to use additional colluding nodes along the tunnel path as application layer waypoints.

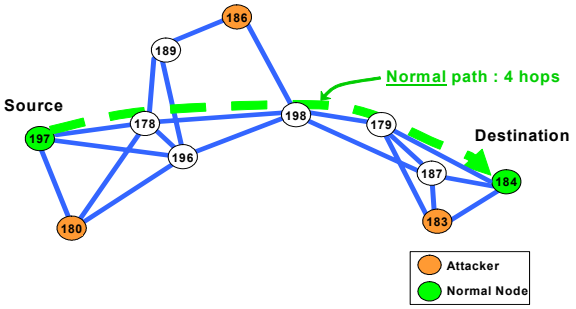

Figure 1. Normal path between nodes 197 and 184

Figure 1 shows the normal route between a source of traffic, node 197, and a destination, node 184. In this example, the route passes through intermediate nodes 178, 198, and 179 and is 4 hops in length. Also shown are three colluding attackers, nodes 180, 186, and 183. Note that the normal path between 197 and 184 does not pass through any of the attacking nodes.

Figure 2 shows the same topology and illustrates the impact of a self-contained in-band wormhole on this path. The wormhole link (the illusory shortcut) stretches between nodes 180 and 183. The tunnel used to create the illusion passes the following sequence of nodes: (180, 178, 189, 186, 198, 187, 183). The role of

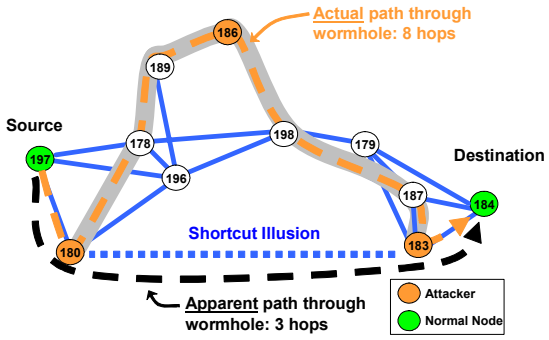

Figure 2. Apparent vs. actual paths during a self-contained in-band wormhole attack

node 186 in the attack is to serve as an application-layer waypoint to prevent tunnel collapse but it also lengthens the tunnel path. Nodes 178, 189, 198, and 187 are unwitting participants.

Once the wormhole is activated, nodes in the network are tricked into believing that a shorter, 3-hop path exists between nodes 197 and 184 via the wormhole link (197, 180, 183, 184). In contrast to the normal path, this apparently shorter path passes through all three attacking nodes, providing each of those nodes an opportunity to control or disrupt traffic between 197 and 184, as well as traffic between other nearby pairs of nodes. As a result, node 197 will forward to node 180 all traffic destined for 184. Node 180 will forward the traffic through the covert tunnel, via node 186, to node 183. When the traffic emerges from the tunnel, node 183 will forward it to 184, its final destination. Traffic from 197 to 184, which appears to be traveling only 3 hops, will in fact travel 8 hops, and pass through each of the attackers.

4. Resilient detection of in-band wormholes

In this section, we discuss our strategies for detecting and localizing in-band wormholes and for making the detector resilient to injection of bogus detection data. We also describe our prototype detection system.

4.1. Detection strategy

The disparity between the apparent path length and the actual path length provides a potential opportunity to detect such attacks. Loss and delay are inherent to the wireless links from which MANETs are constructed. Since loss and delay are cumulative, the wormhole link, which is actually is a 6-hop tunnel, is likely to exhibit substantially higher loss and delay characteristics than a true, single-hop link. Our detection strategy is based on this premise, namely,

that wormholes will have a measurable impact in terms of packet timing or loss on the traffic they carry. This premise has been investigated by other researchers [10]. A wormhole tunnel that provides reliable end-to-end transport can completely mask the cumulative effects of loss, which, for long tunnels, may be so pronounced that OLSR refuses to recognize the wormhole link's existence [14]. However, a reliable tunnel will transform path loss into packet delays, because lost packets will not be retransmitted until acknowledgements have timed out. Since our test topologies involve long tunnels, our tests will use reliable wormhole tunnels and delay-based detection. More specifically, our detection sensors will designate each path measured as either anomalous or normal, depending on the amount of roundtrip delay that is observed.

The simplest way to detect the presence and location of the self-contained wormhole link shown in Figure 2 would be to obtain roundtrip loss and delay statistics from nodes 180 and 183. These nodes, however, are the attackers. One cannot rely on them to report such statistics correctly, as doing so would mean voluntarily identifying themselves to the intrusion detection system. If the attackers are sophisticated, they may know how the intrusion detection system works and may lie by reporting loss and delay statistics that are typical for a true link. This illustrates how a cooperative intrusion detection system can be undermined by injecting bogus data.

To help focus our research on this specific problem, we make the following assumptions:

- There are at most two colluding attackers in any 3-hop neighborhood.
- Probe packets measuring roundtrip loss and delay can be made stealthy so that intermediate nodes along the path can neither distinguish them from other packets nor give them preferential treatment.
- Measurements can be communicated reliably to one or more correlation nodes. (This assumption is discussed further in Section 7.)
- Hop-by-hop and end-to-end authentication services, together with recent knowledge of the local topology, can be used to prove that a measurement probe followed the path claimed by the reporting node.

4.2. Byzantine resilience strategies

In this section, we present three strategies for making the detection of self-contained in-band wormholes resilient to Byzantine behavior by detection sensors.

4.2.1. "No self-reports". As discussed above, if nodes on which detection sensors reside may be compromised, it makes no sense to collect loss or delay measurements about link characteristics from the nodes that control those links; doing so amounts to asking nodes to report on their own misbehavior. Instead, our first strategy is to rely only on "third-party" observations. Instead of collecting link measurements directly, we collect measurements of longer paths that pass through the link, obtaining these from neighbors of the link's endpoints. We then attempt to infer the loss or delay characteristics of the link by correlating multiple measurements. This is illustrated in Figure 3.

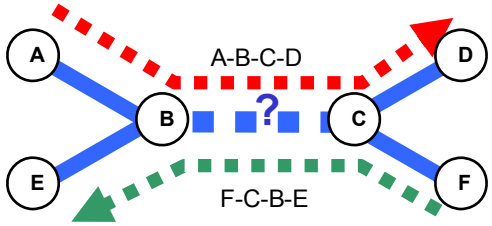

Figure 3. Using 3rd party observations to infer characteristics of link BC

To determine the characteristics of link BC, the neighbors of B and C collect roundtrip measurements of the 3-hop paths that pass through BC. In this example, node A collects roundtrip measurements of path ABCD and node F collects roundtrip measurements of FCBE. If loss or delay measurements along those two paths are anomalous, and if there is at most one wormhole link in this neighborhood, then it must lie at the intersection of these paths, the link BC. In this way, nodes A, D, E, and F are acting as observers of link BC, and of the nodes that control it: B and C. This avoids relying on B and C to report on themselves.

Note that a 3-hop path is the *shortest* path that avoids this problem. Using measurements from longer paths might be workable, but increases the difficulty of determining which link along the path is responsible for the measured anomaly. In general, making this determination would require correlating observations from a larger pool of observers.

4.2.2. Require independent corroboration. An anomalous measurement along a single 3-hop path may indicate that a wormhole is present, but cannot identify which of the three hops is the wormhole link. To localize the wormhole, two measurements through the wormhole link must be made by independent pairs of observers; in other words, the paths must have disjoint endpoints. As shown in Figure 4, if two measurement

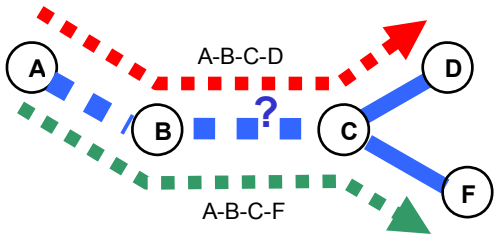

Figure 4. Unless path measurements come from independent observers, they cannot reliably localize the wormhole link

paths have the same source (node A) and traverse the same middle hop, then their intersection includes two links (AB and BC), so they cannot localize the wormhole link. Furthermore, they cannot be assumed to be independent, because the common source (node A) could fabricate measurements for both paths. Similarly, if two paths share a common destination instead of source, independence still cannot be assumed, because that destination can perturb both sets of measurements. Only by having disjoint endpoints can measurements be treated as coming from independent pairs of observers. This leads to our second strategy: detecting a self-contained in-band wormhole should require receiving corroborating anomalous path measurements from at least two independent pairs of observers.

4.2.3. Voting – majority rules. By requiring independent corroboration, the above strategy prevents a single malicious node from being able to launch a successful false accusation against a victim link. However, two colluding attackers that are immediate neighbors of a victim link may be able to. For example, in Figure 3, if any two of the nodes A, D, E, and F are attackers, they may be able to corroborate each other's false accusations of link BC. One way to avoid this possibility would be to require more than two pairs of independent observers. The drawback is that some true wormholes might escape detection for lack of sufficient neighbors.

An alternative strategy is to weigh the balance of positive and negative votes before determining whether a link is anomalous or normal, i.e., seek a consensus of local observers. As shown in Figure 5, if a sufficient number of honest neighbors report normal measurements through link BC, then the preponderance of evidence is that the accused nodes are innocent. This strategy has the potential for counteracting false accusations for more than two colluding attackers, but only if at least as many independent pairs report normal measurements as report anomalous measurements. Otherwise the

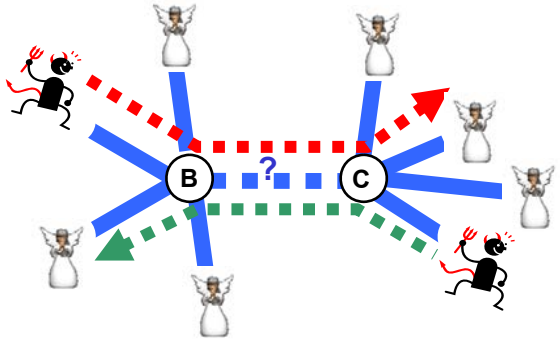

Figure 5. Weighing the balance of positive and negative evidence can counteract corroborating false accusations

attackers will outvote the honest nodes. As discussed below, preventing this from happening requires the presence of more honest observers than attackers.

4.3. Vote tallying and false accusations.

Figure 6 shows examples of how votes are tallied and how independent pairs of observers are counted. In the figure, nodes V1 and V2 are falsely accused victims, A1 and A2 are attackers, and the remaining nodes (B through G) are honest observers. We assume here that if an attacker is either the originator of a roundtrip measurement or the turnaround point, the measurement will be reported as anomalous. In other words, if a pair of observers is either malicious or mixed, an anomalous vote will result. Only if both observers are honest will a normal measurement result. This asymmetry means that a tie vote will require twice as many honest nodes as malicious nodes.

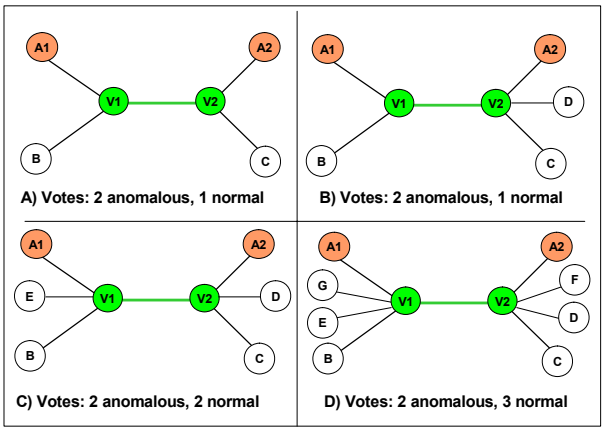

Figure 6. Vote tally examples

In a previous section, we stated that only by having disjoint endpoints can measurements be treated as coming from independent pairs of observers. This is a

slight oversimplification. In actuality, we apply this vote tallying rule only if the measurements in question are in agreement. If two measurements from the same source disagree (if one is normal and the other anomalous), we make an exception and treat them as coming from independent pairs. This rule is necessary to prevent attackers from disqualifying honest nodes from voting with other honest nodes.

In Figure 6, consider example 6A. Two pairs of independent observers can falsely report anomalous votes, namely (A1,C) and (A2,B). The pair (B, C) should be able to report a single normal vote. B however, is also participating in measurements with A2. Similarly, C is also participating in measurements with A1. So without this exception, all normal votes might be disqualified; this gives the attackers an unfair and unnecessary advantage because the pair (B,C) should be entitled to register one normal vote.

In example 6B, an additional honest node D is added, but this does not increase the number of honest votes because D requires an additional honest node to pair up with. In example 6C, honest node E is added, which makes another independent pair of observers possible. Note that D could be paired with either B or E, as may C. Regardless, two independent pairs of honest observers are now possible, which brings the tally to a 2-2 tie. As example 6D shows, to reach a 2-3 tally in favor of the honest nodes, at least 6 honest nodes are needed, with at least 3 on each side of the accused link.

4.4. Self-contained wormhole detection prototype

The strategies presented above are intended to support accurate wormhole detection despite injection of bogus intrusion detection data by colluding attackers. This means being able to tolerate false accusations while ensuring that false claims of innocence do not undermine correct detection and localization of attacking nodes. To assess the utility of these strategies, we constructed a prototype distributed detection system and subjected it to a variety of test conditions in a 48-node network testbed running the OLSR protocol.

As in our prior work [14], the prototype is a distributed cooperative detection system in which each node in the network periodically measures the roundtrip delay to each of the nodes that OLSR states are 3 hops away. These measurements are obtained by using *ping –R*. Each node measures these roundtrip delays using 2 ping packets every 5 seconds plus random jitter between 0 and 5 seconds. The –R option records the roundtrip path taken by the probes. If a ping packet ends up taking a different path toward the destination than returning, the associated delay measurements are discarded. As stated in Section 4.1, we assume that probe packets, unlike ping packets, can be made stealthy. Consequently, we use ping simply as a placeholder for a more-sophisticated future probe mechanism.

Nodes send anomalous and normal delay measurements and path descriptions to a correlation node that is the root of a dynamic hierarchy [23]. The correlation node attempts to determine whether a wormhole is present, and if so, the location of the wormhole link. Paths whose minimum delay times are greater than a specified threshold are classified as anomalous. To designate a link as a wormhole link, the correlator requires the following evidence:

- anomalous measurements from at least two independent pairs of observers; and
- fewer normal measurements from independent paths through the same (accused) middle link than anomalous measurements.

In case of a tie between anomalous and normal votes, the benefit of the doubt goes to the accused link, which will not be identified as a wormhole.

5. Testing

Below we describe our testbeds, attack tools, and test scenario.

5.1. Network testbeds

Our test environment consists of two testbeds based on NRL's Mobile Ad-Hoc Network Emulator (MANE) [19]. One testbed is at ARL; the other is at SPARTA. Both contain 48 test nodes; one or more MANE servers, which emulate node positions, mobility, and radio connectivity; and a single experiment control and monitoring node. Each test node is connected to a single MANE server. Using geographic position data periodically advertised by the controller node, each MANE server computes the current MANET connectivity, and selectively forwards, corrupts, or drops packets traveling between test nodes. MANE servers share a separate server-to-server data channel over which packets are forwarded between test nodes residing on different servers. Test nodes are also connected to the experiment control node via a separate channel used for control, monitoring, and data collection. Each test node runs the Fedora Core 3 operating system and the OLSR daemon developed at the University of Oslo's UniK organization [22].

ARL's testbed includes four MANE servers, each connected to twelve test nodes. Each test node in the

ARL testbed contains a Pentium 4 3GHz processor, 1 GB of RAM, and two gigabit Ethernet interfaces, which are used to connect the node to its assigned MANE server and to the control channel.

SPARTA's testbed includes a single MANE server, which is connected to all 48 test nodes via VLAN-configured Ethernet switches. The test nodes are a heterogeneous collection of older Pentium-family systems having CPU speeds from approximately 400 MHz to 2 GHz, and memory capacities from 128MB to 1GB. The test nodes are connected to the MANE server and control channel via 100 mbps Ethernet links.

5.2. Attack tools

We used two types of attack tools. The first, described in more detail previously [14], creates in-band wormholes. This requires creating multiple tunnels, capturing OLSR control messages, forwarding them through the tunnels, and, for some forms of wormholes, broadcasting the messages after they emerge from the tunnels. After the wormhole begins attracting data packets, these too are forwarded through the tunnels. In addition, the routes on each attacking node must be specially configured according to the surrounding network topology to ensure that tunneled data packets are correctly forwarded.

The second type of tool causes bogus information to be inserted into the intrusion detection system, directly or indirectly. It includes two components. One is modified version of the *data collector*, an element of the intrusion detection system that runs on all nodes. The benign version of the data collector sends roundtrip probes to 3-hop neighbors. When a probe returns, the data collector reports the measured loss and delay to the intrusion detection system. Attacking nodes use a malicious version of this component to report false values to the intrusion detection system to implicate other nodes or evade detection. The other component uses the Linux *ipqueue* subsystem [18] to delay roundtrip probes it receives and returns. This corrupts the measurements observed by an honest probe initiator, causing the initiator to report the associated link as anomalous. These false accusations and measurement delays are activated by these components only when two attackers are three hops apart and are potentially able to create corroborating false accusations.

5.3. Test scenario

An attacker's ability to attract traffic and consume bandwidth during a wormhole attack depends on the positioning of the attacking nodes within the MANET topology. For example, if the attacking nodes are all relatively close together, the wormhole link may not appear to be a significant shortcut for any traffic sources and destinations. We developed a specialized mobility scenario that is conducive to usable wormhole attacks yet provides sufficient topological randomness to test the prototype detector under a variety of conditions. In particular, the test scenario was designed to maximize the opportunity for attackers to generate corroborating false accusations to test our opportunistic voting algorithm.

In this scenario, benign nodes are assigned random positions in a two-dimensional field that is 500 meters in width and 350 meters in "height". These nodes travel at a speed of 1 meter per second along randomly assigned headings. When a node reaches the field boundary, it is reflected and continues onward. The mobility scenario includes four colluding attackers. The attacking nodes are initially placed at the corners of a 450 meter by 250 meter inner rectangle that is centered within the field. Attacking nodes also travel at 1 meter per second along random headings; however, each attacker is additionally confined to a small rectangular region 15 meters by 15 meters. This allows some attacker movement, but maintains the approximate initial positioning near the corners of the inner rectangle. Figure 7 shows an example topology.

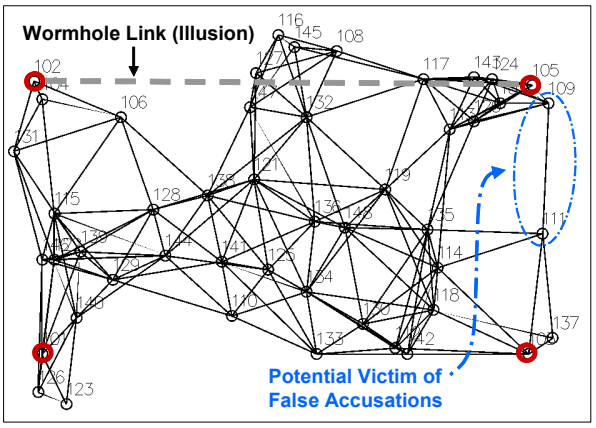

Figure 7. Example 48-node topology

This attacker positioning satisfies two objectives. First, the wide horizontal spacing allows the upper (or lower) pair of attackers to create a wormhole link that appears to be substantially shorter than any real path from one side of the topology to the other. Such a wormhole link will exhibit a strong "gravitational" force [11], attracting traffic from large regions around the wormhole endpoints. Second, the vertical positioning was chosen so that the pair of attacking nodes on each side of the inner rectangle would often be 3 hops apart, given the 120-meter emulated

transmission range we chose. This enables each such pair to generate corroborating false accusations, as described in Section 4.2 above.

As shown in Figure 7, we configured the wormhole tool to create a self-contained wormhole link (shortcut) between the top left and top right attackers, nodes 102 and 105. The covert tunnel used to carry OLSR control packets and data packets from 102 to 105 traverses the other three sides of the attacker rectangle, traveling via attacker waypoints 101 (bottom left) and 103 (bottom right). Since the attacker nodes are not immediate neighbors, the tunnel is necessarily routed through benign intermediary nodes, who serve as unwitting accomplices. Although affected by geographic positioning, the number of hops between a pair of attackers also depends on the positioning of these benign intermediaries.

This wormhole configuration results in a long tunnel, often 9 or more hops in length. From the attacker's perspective, a long tunnel may be desirable because it consumes more bandwidth and does so across more neighborhoods. Hence this attack is relatively severe. On the other hand, a wormhole with a long tunnel should be easier to detect because the tunnel path's loss and delay characteristics should be clearly distinguishable from the characteristics of a true single hop, which the tunnel purports to be. In fact, this configuration was chosen to facilitate sensor accuracy so that our testing could focus on the problem of interest: countering deliberate lying by malicious sensors rather than mitigating honest errors by benign sensors. Figure 7 also shows that the link between nodes 109 and 111 is a potential victim of corroborating false accusations, because it lies in the middle of a 3-hop path between two attacking nodes, 105 and 103.

In order for OSLR to conclude that nodes 102 and 105 are immediate neighbors and can be used to improve routes, OLSR control packets (e.g., HELLOs) need to travel reliably from one end of the tunnel to the other. In the presence of probabilistic packet loss of even a few percent, a UDP-based tunnel of this length will not deliver HELLOs with sufficient consistency to satisfy OLSR, and the wormhole link will not be persistent. To avoid this problem, the wormhole tunnel was configured to provide a reliable transport based on TCP. Consequently, we configured the prototype to use measurements of roundtrip delay rather than loss. This is based on the expectation that the acknowledgement timeouts and packet retransmits triggered by probabilistic loss will exaggerate roundtrip times through the tunnel.

The attacker software is configured to inject three kinds of bogus information into the intrusion detection system. First, when reporting 3-hop roundtrip delay times that pass through wormhole link (e.g., as the first hop), the attacker will falsely report that roundtrip times are normal. The intent of this lie is to evade detection. In reality, because the detector does not consider "self-reports", this form of lying should have no impact on the detector. Second, whenever an attacker is three hops from its counterpart on the same side of the topology, it will report that roundtrip times to each of its 3-hop neighbors are anomalous. Third, under the same conditions, the attacker will delay 3-hop probes for which it is the roundtrip turnaround point. If such a probe was generated by an honest node, the node will be tricked into reporting the path as anomalous. If 3-hop probes *to* or *from* a pair of attackers traverse a common middle hop, the two attackers will then be able to generate (or stimulate) corroborating false accusations. If a sufficient number of honest nodes are not topologically positioned to probe this victim link and refute the accusation, the attackers will then be able to trick the detector.

6. Test results

We conducted two sets of tests. The objective of the first set was to verify that the prototype detector is able to detect self-contained in-band wormholes, despite attempts at evasion by the attacker. These tests were run on ARL's testbed because its faster, more reliable hardware improved the stability and duration of the wormholes, providing more usable detection testing opportunities per run. The objective of the second set was to assess the resilience of the detector to corroborating false accusations by colluding attacker pairs. Since these tests do not require persistent stable wormholes, we were able to run them on SPARTA's testbed.

6.1. Wormhole detection test results

The first test consisted of 60 cycles, each containing 10 seconds of topology changes generated by the random heading mobility scenario described above, followed by a wormhole attack lasting 90 seconds. The topology was held constant during the wormhole attack. Each cycle included time for OLSR routes to settle – 20 seconds before and 30 seconds after the attack. The testbed was also configured to emulate packet loss based on a free-space path loss model. Although the network carried OLSR control traffic, wormhole tunnel control traffic, and intrusion detection infrastructure traffic, no application traffic load was present. The detector's path delay threshold for designating a 3-hop path as anomalous was set to 6 ms. This value was determined via experimentation.

To constitute a valid wormhole detection test case, we require a wormhole to have a total lifetime exceeding 15 seconds. Of the 60 test cycles, 47 resulted in wormholes that met this criterion, including several intermittent wormholes that existed over multiple time intervals, each shorter than 15 seconds.

Despite the attackers' attempts to evade detection by reporting normal roundtrip times for paths traversing each wormhole, the prototype successfully detected all 47 valid wormholes. In addition, it detected four other wormholes having total lifetimes shorter than 15 seconds. The detector also generated one false alarm, in which it erroneously identified a normal link as a wormhole. This test was not intended to assess the prototype's accuracy with any depth or precision, but simply to demonstrate that detection of wormholes by measuring and analyzing roundtrip delays is feasible under certain circumstances.

6.2. False accusation test results

The second set of tests consisted of twelve, hour-long test runs. Each run used 48, 43, 38, or 33 nodes, resulting in four different network densities. The 43-node configuration was created by discarding 5 random nodes from the 48-node configuration. The 38-node configuration was created by discarding 5 random nodes from the 43-node configuration, and so on. At each density, three runs were conducted using different random seeds (A, B, or C) to determine the initial headings of nodes. All runs used the attacker positioning and mobility scenario described in Section 5.3. However, unlike the wormhole detection test, the false accusation tests utilized continuous node mobility.

We introduce the following terminology to explain the false accusation test results. A corroborated false accusation or "FA" is a false accusation against a victim link comprising anomalous measurements reported by at least two independent pairs of observers. A false accusation reported by a single pair of observers, but not corroborated by other independent pairs, is of little interest here because the prototype detector has been designed to ignore uncorroborated accusations.

We further distinguish between two kinds of (corroborated) FAs. A successful FA is an FA that *is not outvoted* by an equal or greater number of normal measurement reports from independent pairs of honest nodes, i.e., the attackers prevail. A suppressed FA is an FA that *is outvoted* by an equal or greater number of normal measurement reports from independent pairs of honest nodes, i.e., the honest nodes prevail.

Table 1. Summary of one-hour test runs

No. of Nodes	Initial Random Seed	Ave Degree	Link-Sec Corrob. FAs	Links Subject to Corrob. FAs	Ave FA-Sec per Link
48	A	7.42	1,344	45	29.87
48	B	7.20	1,663	59	28.19
48	C	7.26	2,202	69	31.91
43	A	6.62	1,106	40	27.65
43	B	6.33	1,235	43	28.72
43	C	6.46	1,203	38	31.66
38	A	5.88	775	32	24.22
38	B	5.51	810	29	27.93
38	C	5.68	1,167	40	29.18
33	A	4.98	583	23	25.35
33	B	4.75	695	28	24.82
33	C	4.93	698	27	25.85
Total			13,481	473	28.50

Table 1 provides a summary of the twelve false accusation test runs. The average node degree (number of one-hop neighbors per node) ranged from 4.75 to 7.42. The number of seconds during which links were subjected to corroborated false accusations ranged from 583 to 2,202 link-seconds. The number of links that were subjected to corroborated false accusations ranged from 23 to 69. As discussed below, these last two statistics indicate that the amount of "interesting" FA data generated during these runs varied and was strongly correlated with density. In total, these runs generated 13,481 link-seconds of FAs against 473 victim links, a significant volume of interesting data.

Figure 8 shows the overall effectiveness of opportunistic voting in instantaneously suppressing FAs across the test runs. We define effectiveness as the percentage of link-seconds during which the FAs were suppressed, i.e., were outvoted by pairs of honest observers. Suppression rates ranged from 57% (run 33B) to 81% (run 43A). The overall weighted average across all runs was 72%, with weights reflecting the amount of FA data generated during each run.

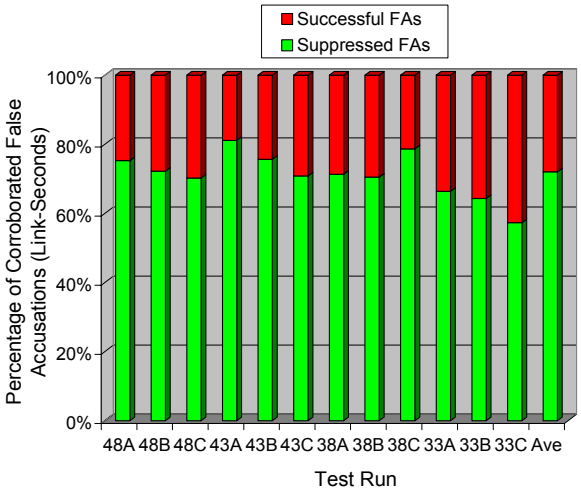

Figure 8. Successful vs. Suppressed False Accusations

Figure 9 depicts the relationship between FA suppression effectiveness and node degree, using averages for each of the four densities. Figure 9 shows that effectiveness was only loosely correlated with average node degree. Not surprisingly, both average node degree (4.89) and effectiveness (62.67%) were lowest during the 33-node runs. However, despite having the highest average node degree, the 48-node runs exhibited effectiveness (72.58%) *lower* than the 43-node and 38-node runs (75.85% and 73.50%). We discuss possible causes for this result in Section 7. Similarly, as shown in Figure 8, the effectiveness of the best 48-node run (48A), was lower than that of the best 43 (43A) and 38 (38C) node runs.

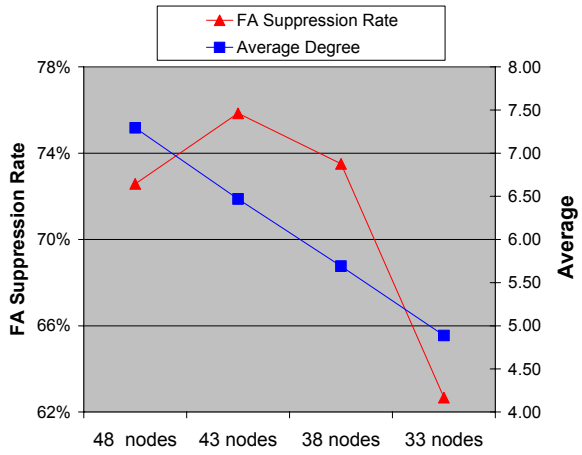

Figure 9. FA Suppression rate and average degree

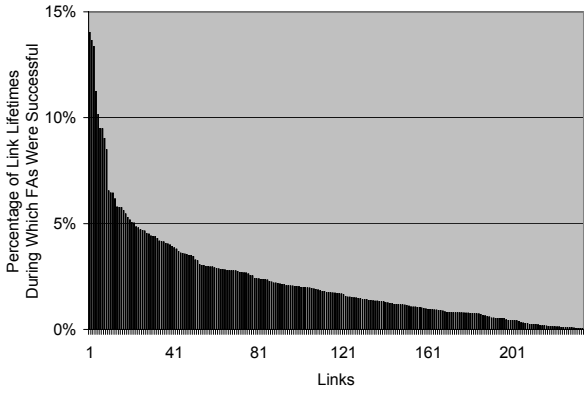

Figure 10. Successful FAs - lifetime percentage

As shown in Table 1 above, across all test runs, 473 links were subjected to FAs. Of these, 234 links were subjected to *successful* FAs. (For the other 239, all FAs were suppressed by honest outvoting.) Figure 10 depicts a per-link lifetime analysis of the 234 successfully accused links. It shows that no link was successfully accused for more than 14% of its lifetime. Conversely for at least 86% of the lifetime of every link, the link was either not subjected to corroborating accusations (because the attackers were not appropriately positioned topologically), or the accusations were suppressed by honest outvoting. The vast majority (212) of the 234 links were successfully accused for less than 5% of their lifetimes.

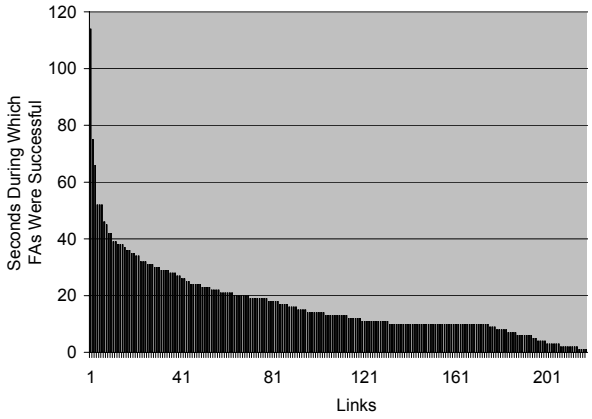

Figure 11. Successful FAs - duration

Figure 11 is a similar graph but shows the duration of successful FAs for these same links. No link was successfully accused for more than a total of 114 seconds. Furthermore, only 13% of these links were successfully accused for more than 30 seconds, and only 4% for more than 40 seconds.

7. Discussion

The 72% overall effectiveness of opportunistic voting in instantaneously suppressing FAs is encouraging, but by itself, is probably not sufficient to allow use in an operational setting. Significantly greater accuracy, however, appears possible by combining voting tallies with other historical (non-instantaneous) evidence. First, since many of the successful FA periods were relatively short, an operational detector could be configured to ignore FA periods that persist for less than a minimum duration, for example, 30 seconds. While this would increase detection latency from the onset of a true wormhole attack, it could be justified if wormhole attacks of such short duration were deemed 1) to pose a minor threat and 2) too difficult to distinguish from the general MANET "noise". Second, since a normal link is unlikely to transform itself into a wormhole (or vice versa), a link's prior or subsequent history can also be considered. If 95% of a link's lifetime has been free of successful accusations, then the 5% during which the link has been accused should be evaluated with skepticism. Third, it may be possible to utilize voting history. Nodes whose accusatory votes have often been at odds with the majority could be given less weight in voting tallies. This would force attackers to be much more selective in use of false accusations because failure to prevail in initial voting tallies would reduce chances of prevailing in the future.

A surprising result is that the effectiveness of opportunistic voting in suppressing FAs does not appear to be more strongly correlated with average node degree. This is counter-intuitive because having a larger population of honest nodes in the neighborhood surrounding a falsely accused link would seem to increase the likelihood that two colluding attackers would be outvoted.

Whether there are enough independent pairs of observers to vote for the accused link depends not only on the density of the network, but on two other factors. First, the topology of the network must be such that the accused nodes have a sufficient number of *disjoint* honest neighbors. For example, if node A in Figure 3 were to move to the right and become a neighbor of both B and C, it would no longer be able to send 3-hop probes to D via link BC. D would become two hops from A and would cease being a destination for A's 3-hop probes. Moreover, even if A sent probes to D, they would pass only through C, skipping B. So despite A's proximity, it would no longer be able to probe or vote for link BC.

Second, between a pair of nodes that are 3 hops apart, there may be several alternate paths available as routes, and OLSR will arbitrarily choose any one of these. If OLSR chooses the path through the falsely accused link, the pair of nodes can probe the path and then vote in favor of the accused link. However, if OLSR chooses a different path, the pair will be able to vote on some other link's innocence, but not the falsely accused link. Although increasing the density in the vicinity of a victim link may increase the potential number of paths through which observers may probe an accused link, the actual number is determined by OLSR.

Since the effectiveness of an in-band wormhole attack and false accusations depends on attacker spacing, it is challenging to construct good test cases that utilize random mobility models. As mentioned above, the amount of interesting FA data generated during test runs was strongly correlated with network density. This probably resulted from using identical geographic spacing for the attackers in all runs. The spacing was chosen so that the pair of attackers on each side of the topology would be 3 hops apart during most of the 48-node test runs. We decided to keep this spacing the same in all runs to minimize the number of experimental variables that might affect run outcomes. However, the number of hops between two nodes depends not only on geographic spacing but on the proximity of intermediary nodes. At lower network densities, there is less chance that two intermediary nodes will be close to the "line-of-sight" between a pair of attackers. If the two closest intermediaries are too far from this line then they will not be able to bridge a path between two attackers without additional intermediaries, which would increase the hop count between the attackers. Consequently, at lower network densities, the length of a path between a pair of attackers is more likely to be greater than 3 hops. The result is that a pair of attackers will less often be able to generate corroborating false accusations. In hindsight, it might have been better to adjust (reduce) the attacker spacing for lower densities to compensate. Alternatively, run durations could be extended for lower density test runs.

A simplifying assumption listed in Section 4.1 is that measurements can be reliably communicated to one or more correlation nodes. In making this assumption, we sidestep one of the fundamental issues underlying the Byzantine agreement problem in distributed systems, namely that communications between nodes often depends on correct forwarding by intermediary nodes, which may be malicious [15]. We focus instead on a prerequisite question:

> *Is the totality of data distributed throughout the network adequate to enable a correct consensus?*

If the data is not adequate, then no communications protocol can make it so, and the forwarding issue is

moot. This paper provides partial evidence that the data produced by our prototype (3-hop path measurements) is adequate for correct wormhole detection decisions, even in the presence of colluding attackers and false accusations. However, reducing these techniques to practice will ultimately require being able to reach consensus when communications depends on forwarding by potentially malicious intermediary nodes.

8. Conclusion

Detecting attacks on MANET routing and forwarding requires cooperative detection techniques that utilize ordinary, relatively vulnerable hosts as intrusion sensors. If compromised, these hosts can inject bogus data into the intrusion detection system to conceal malicious activities or falsely accuse well-behaved nodes. The ease with which such data can be injected poses a fundamental threat to any cooperative intrusion detection system. Approaches to Byzantine fault tolerance that involve voting are potentially applicable, but must address the fact that only nodes in particular topological locations at a particular time 1) can observe the symptoms of a particular attack and 2) are therefore eligible to vote on whether an attack occurred.

We have examined this issue in the context of a prototype distributed detector for self-contained, in-band wormholes in an OLSR network, focusing in particular on false accusations. The detector analyzes roundtrip delay measurements of 3-hop paths and correlates these to identify the locations of wormhole links and thus the attackers. The detector was tested using specialized wormhole attack software that reports bogus delay measurements to conceal its presence and falsely accuse nearby nodes. In addition, this software manipulates delay measurements reported by other nodes so that they are tricked into falsely accusing their neighbors.

We tested the prototype detector in a 48-node wireless emulation testbed employing random node mobility and attacker spacing conducive to corroboration of false accusations. We conducted 12 hour-long tests using four network densities and three random initialization seeds. The results show that during 13,481 link-seconds of corroborated false accusations, opportunistic voting instantaneously suppressed the accusations 72% of the time. Moreover, most periods of successful false accusations were relatively brief, e.g., less than 30 seconds, suggesting that most unsuppressed accusations could be ignored safely at the cost of increased detection latency. Significantly greater accuracy appears achievable by combining instantaneous voting with statistics from each link's prior history; if a link's lifetime has been primarily free of successful accusations, a small period of successful accusations against it should be treated with skepticism. For the network densities tested, the effectiveness of voting in suppressing false accusations was only loosely correlated with average node degree.

The work described here has focused on tolerating false accusations. Our future work will include treating false accusations as yet another intrusion symptom that can be used to identify attackers. This may be complex because attackers can trick other nodes into generating false accusations and may do so probabilistically to remain stealthy. We have recently begun using game theoretic techniques to analyze such interplay between wormhole countermeasures and counter-countermeasures [6]. These techniques should enable the formulation of mathematically-based strategies that optimize intrusion detection and response system effectiveness.

9. References

[1] B. Awerbuch, R. Curtmola, D. Holmer, C. Nita-Rotaru, and H. Rubens, "Mitigating Byzantine Attacks in Ad Hoc Wireless Networks", Technical Report Version 1March 2004, Department of Computer Science, Johns Hopkins University.

[2] B. Awerbuch, R. Curtmola, D. Holmer, H. Rubens, C. Nita-Rotaru, "On the Survivability of Routing Protocols in Ad HocWireless Networks", SecureComm 2005 - First International Conference on Security and Privacy for Emerging Areas in Communication Networks, September 2005.

[3] C. Adjih, T. Clausen, A. Laouiti, P. Muhlethaler, and D. Raffo, "Securing the OLSR routing protocol with or without compromised nodes in the network", INRIA, Tech. Rep. ISRN INRIAR/RR-5494, February 2005.

[4] Buchegger, S.; Le Boudec, J.-Y., "Nodes Bearing Grudges: Towards Routing Security, Fairness, and Robustness in Mobile Ad Hoc Networks", Proc. 10th Euromicro Workshop on Parallel, Distributed and Network-based Processing, pp. 403-410, Canary Islands, Spain, January 2002.

[5] S. Buchegger, J. Le Boudec, "Coping with False Accusations in Misbehavior Reputation Systems for Mobile Ad-hoc Networks", Technical Report IC/2003/31, EPFL-DI-ICA (2003).

[6] J. Baras, S. Radosavac, G. Theodorakopoulos, D. Sterne, P. Budulas and R. Gopaul, "Intrusion Detection System Resiliency to Byzantine Attacks: The Case Study of Wormholes in OLSR", MILCOM 2007, Orlando, FL, October 2007.

[7] L. Buttyán, L. Dóra, and I. Vajda, **"**Statistical Wormhole Detection in Sensor Networks", Second

European Workshop on Security and Privacy in Ad Hoc and Sensor Networks (ESAS 2005) Visegrád, Hungary, July 13-14, 2005.
[8] R. Castro, M. Coates, G. Liang, R. Nowak and B.Yu. "Network Tomography: Recent Developments", *Statistical Science* 2004, Vol. 19, No. 3, 499–517.
[9] Y. Hu, A. Perrig, and D. Johnson, "Packet Leashes: A Defense against Wormhole Attacks in Wireless Ad Hoc Networks", Proceedings of IEEE Infocomm 2003.
[10] M. Gorlatova, P. Mason, M. Wang, L. Lamont, R. Liscano, "Detecting Wormhole Attacks in Mobile Ad Hoc Networks through Protocol Breaking and Packet Timing Analysis", MILCOM 2006, October 2006, Washington, DC.
[11] R. Gopaul, P. Kruus, D. Sterne, B. Rivera, "Gravitational Analysis of the In-Band Wormhole Phenomenon", Proc. 25th Army Science Conference, November 27-30, Orlando, FL.
[12] Islam, M.M. Pose, R. Kopp, C. "An Intrusion Detection System for Suburban Ad-hoc Networks" TENCON 2005 2005 IEEE Region 10, Nov. 2005, Melbourne Australia.
[13] I. Khalil, S. Bagchi, N. B. Shroff, "LITEWORP: A Lightweight Countermeasure for the Wormhole Attack in Multihop Wireless Networks", pp. 612-621, 2005 International Conference on Dependable Systems and Networks (DSN'05), 2005.
[14] P. Kruus, D. Sterne, R. Gopaul, M. Heyman, B. Rivera, P. Budulas, B. Luu, T. Johnson, N. Ivanic, G. Lawler, "In-Band Wormholes and Countermeasures in OLSR Networks", Second International Conference on Security and Privacy in Communication Networks, (SECURECOMM 2006), Aug. 28, 2006, Baltimore, MD.
[15] L. Lamport, R. Shostak, and M. Pease, "The Byzantine Generals Problem", ACM Trans. Programming Languages and Systems, Vol. 4, No. 3, July 1982, pp. 382-401.
[16] L. Lazos, R. Poovendran, C. Meadows, P. Syverson, L. W. Chang, "Preventing Wormhole Attacks on Wireless Ad Hoc Networks: A Graph Theoretic Approach", IEEE Wireless Communications and Networking Conference (WCNC), 2005.
[17] S. Marti, T. J. Giuli, K. Lai, and M. Baker. Mitigating Routing Misbehavior in Mobile Ad Hoc Networks. Proceedings of the 6th Intl. Conference on Mobile Computing and Networking. Boston, MA, August 2000.
[18] *The Netfilter.org Project*, http:///www.netfilter.org .
[19] Mobile Ad-hoc Network Emulator (MANE), http://cs.itd.nrl.navy.mil/work/mane/index.php.
[20] M. Natu, A. Sethi, "Active Probing Approach for Fault Localization in Computer", 4th IEEE/IFIP Workshop on End-to-End Monitoring Techniques and Services, April 2006.
[21] "Optimized Link State Routing (OLSR)", IETF RFC 3626, T. Clausen, P. Jacquet, Ed., October 2003.
[22] *OLSR.org*, http://www.olsr.org.
[23] D. Sterne, P. Balasubramanyam, D. Carman, B. Wilson, R. Talpade, C. Ko, R. Balupari, C-Y Tseng, T. Bowen, K. Levitt, J. Rowe, "A General Cooperative Intrusion Detection Architecture for MANETs", Proc. Third IEEE International Information Assurance Workshop, College Park, MD, March 2005.
[24] Y. Zhang and W. Lee, "*Intrusion Detection in Wireless Ad Hoc Networks*", Proceedings of The Sixth International Conference on Mobile Computing and Networking (MobiCom 2000), Boston, MA, August 2000.
[25] C. Zouridaki, B.L. Mark, M. Hejmo, and R. K. Thomas, "A Quantitative Trust Establishment Framework for Reliable Data Packet Delivery in MANETs", Proceedings of the 3rd ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN 2005), Held in conjunction with the 12th ACM Conference on Computer and Communications Security (CCS 2005), Alexandria, VA, USA, November 7, 2005.
