2011 Eighth IEEE International Conference on Mobile Ad-Hoc and Sensor Systems

Secure Neighbor Discovery in Mobile Ad Hoc Networks

R. Stoleru, H. Wu, H. Chenji
Department of Computer Science and Engineering, Texas A&M University
{stoleru, bylike, cjh}@cse.tamu.edu

***Abstract*—Neighbor discovery is an important part of many protocols for wireless adhoc networks, including localization and routing. When neighbor discovery fails, communications and protocols performance deteriorate. In networks affected by relay attacks, also known as wormholes, the failure may be more subtle. The wormhole may selectively deny or degrade communications. In this paper we present Mobile Secure Neighbor Discovery (MSND), which offers a measure of protection against wormholes by allowing participating mobile nodes to securely determine if they are neighbors. To the best of our knowledge, this work is the first to secure neighbor discovery in mobile adhoc networks. MSND leverages concepts of graph rigidity for wormhole detection. We prove security properties of our protocol, and demonstrate its effectiveness through extensive simulations and a real system evaluation employing Epic motes and iRobot robots.**

I. Introduction

Neighbor discovery is the process by which a node in a network determines the total number and identity of other nodes in its vicinity. It is a fundamental building block of many protocols including localization [1], routing [2], leader election [3], and group management [4]. Time-based communications and many media access control mechanisms [5] rely on accurate neighbor information. Neighbor discovery is especially important to the proper functioning of wireless networks.

In wireless networks, neighbors are usually defined as nodes that lie within radio range of each other. Thus, neighbor discovery can be considered as the exploration of the volume of space or "neighborhood" immediately surrounding a wireless node. Nodes found within the neighborhood are neighbors and, depending on network configuration and topology, may cooperate in the performance of various tasks including communications, sensing and localization. However, wireless communications are susceptible to abuse. Attackers have the freedom to perform malicious activities ranging from simple denial of service to sophisticated deception.

One particularly insidious threat to a wireless network is the wormhole or relay attack [6]. In this attack, two or more attackers collaborate to record communications at the packet or bit level in one location and play them back elsewhere. Wormholes may disrupt communications, alter routing, or induce localization errors. Further exploitation of wormhole-enabled communications can lead to unauthorized physical access, selective dropping of packets, and even denial of service [7]. When a wormhole convinces distant nodes that they are neighbors, it can monitor their traffic as well as any traffic they route across the link. When several nodes in a region are similarly compromised, all communications from the region can be "herded" through narrow chokepoints. At these chokepoints, traffic is consolidated and the attacker can gain maximum knowledge and advantage with minimum effort. The chokepoint also allows the attacker to perform both individually targeted and large scale DoS against any traffic passing through it. Similarly, node localization protocols can be severely impacted (e.g., a node may be mislead to believe that it is within 1 hop of a very distant anchor - a node with known location). Given the potentially severe effects of a wormhole, nodes must be able to securely conduct neighbor discovery.

To address the aforementioned challenges, we present Mobile Secure Neighbor Discovery (MSND), which allows neighbors to verify that they are speaking directly with each other. A wormhole can be detected due to the fact that the path traveled by a ranging signal varies from expected values when a wormhole is present. Instead of traveling directly to the remote node, the ranging signal must travel to one end of the wormhole, transit the wormhole, and then exit to arrive at the destination node. In the case of a static network, this variation is difficult to detect because, for a single node, it is constant. However, node mobility and graph rigidity concepts allow participating nodes to identify distortions caused by wormholes. MSND provides a measure of protection against the threat of wormholes.

The contributions of this paper include:

- A protocol (MSND) for detecting the presence of wormholes when mobile nodes participate.
- Security analysis and correctness of MSND.
- Performance evaluation through simulations, demonstrating accurate wormhole detection with low false negatives.
- A real system evaluation employing Epic motes and iCreate robot hardware, demonstrating the performance of our proposed solution.

II. Related Work

Secure neighbor discovery (SND) covers a range of techniques and technologies. A variety of approaches have been proposed to handle SND in general and wormholes in particular. Many approaches leverage physical properties of communications and can be roughly divided into solutions based on location, time, time and location, and network geometry. Other

978-0-7695-4469-4/11 $26.00 © 2011 IEEE
DOI 10.1109/MASS.2011.15

solutions rely on security properties achievable in specific scenarios. In [7], Papadimitratos, et al. give an overview of the problems and challenges associated with SND. Their paper includes a set of real-world examples illustrating various threats to neighbor discovery.

Location-based solutions offer neighbor discovery protocols to ensure that nodes claiming to be neighbors share the same neighborhood. Coordinated use of both RF and ultrasonic emitters was proposed by Priyantha [8]. Relying on the difference in time of flight between RF and ultrasonic signals, Cricket produces relatively accurate localization both static and mobile nodes at ranges on the order of meters. [9] uses localized beacons to detect wormholes while executing a localization protocol for statically deployed nodes. A mechanism for geographically assigning local broadcast keys was used in [10] to limit the range of communications. However, location-based protocols assume the availability of localization information, at least for a subset of participating nodes, making them unsuitable for scenarios without this information.

Time-based solutions attempt to leverage time-of-flight measurement to ensure that transmitting nodes lie within the local neighborhood. Packet leashes are a well-known example of this approach. Using both geographic and temporal leashes, Hu, et al. [6] propose mechanisms that incorporate high-resolution synchronized clocks to calculate the time or distance of flight of a packet. However, the high level of precision needed exceeds the capabilities of most modern hardware at distances less than kilometers. SECTOR [11] proposed tracking nodes encounters and using these encounters for verification of identity. As the authentication phase of SECTOR relies on nanosecond clocks and special hardware, it is impractical for many adhoc networks. Time-based solutions, however, all face a common constraint. In [12], Poturalski, et al. offer an impossibility proof showing that time-based protocols cannot guarantee SND unless the environment is free of obstacles and the distance between neighbors is small. In mobile adhoc environments, these constraints would require nodes to have constant, detailed location information in order detect and avoid obstacles and node communications would be limited to the period of time when the nodes were close to each other. [12] and [13] offer a general class of alternatives known as time-and-location protocols and provide in-depth theoretical analysis of a solution.

Time-and-location protocols use both time-of-flight measurements and node locations to support SND. Shokri, et al. [14] combine ultrasonic time-of-flight ranging with simple geometric tests to securely verify static neighbors and offer good analysis of the security properties of the protocol. Their solution relies on robust quadrilaterals [15] to establish a geographic relationship between nodes. While effective in static, relatively dense environments, their solution is not viable in low density networks and does not support mobility. Our paper builds on the foundation laid in [14].

Geometry-based solutions detect wormholes by analyzing metrics provided by routing protocols within the network. Xu, et al. [16] use flooding to establish hop distances between nodes. The resulting map is analyzed to detect wormholes. Maheshwari, et al. [17] assemble a local connectivity graph and analyze it for forbidden structures created by the wormhole. Both of these papers assume a relatively high degree of connectivity. Additionally, reliable hop metrics can be prohibitively costly to maintain in networks with mobility. Other solutions use a centralized approach to create rigid graphs [18], or statistically measure wormhole-induced distortion of the average number of neighbors and average shortest path lengths throughout the entire network [19]. Both of these centralized solutions assume continuous connectivity and the presence of a sink or network controller.

A final set of approaches to SND relies on properties achievable only in certain contexts. Liu [20] describes SND as a problem of neighbor validation and assumes that attacker capabilities are limited during initial sensor deployment. Nodes securely determine neighbors during this period. Validation is handled through neighbor table exchanges and requires a static and well-connected network. Directional antennas were proposed as a defense against wormholes in [21]. Although effective, the addition of this type of hardware is limiting and costly in many wireless network deployments.

III. Preliminaries and Problem Formulation

Wormholes pose a subtle, insidious threat because they can affect communications without directly participating as network entities. However, wormholes introduce observable changes in a network. When nodes are static, this variation is difficult to detect. For example, in [14], the solution requires four or more favorably positioned static nodes to accurately measure wormhole-induced changes. This section presents the system and threat models, and problem formulation.

A. System Model

Our system model is motivated by DistressNet [22], a wireless sensor, adhoc and delay tolerant network system for disaster response and military applications. DistressNet employs wireless adhoc and delay tolerant networks, consisting of wireless mobile nodes (e.g., emergency or military vehicles, mobile equipment, emergency responders, military personnel, etc.) distributed across a 2D region. For DistressNet applications, we assume that not all nodes have GPS and that the environment is GPS-denied (such as in military). Consequently, a secure neighbor discovery protocol becomes essential for wireless mobile nodes to correctly obtain their location.

Each node is equipped with a single radio transceiver, a ranging capability, and a clock with enough precision to support ranging operations (e.g., hundreds of microseconds precision for 0.5-1.5m ranging accuracy, for acoustic/ultra-sonic ranging). Communications between nodes use bidirectionally symmetric radio transmissions with a range R_{RF}. Ranging radius, R_{RNG}, is similarly bidirectional and symmetric. Nodes are real neighbors if they can communicate via radio and perform ranging operations with each other. Mobile nodes are able to calculate distance traveled with some degree of error

(e.g., 2%-10% of the distance traveled, using dead-reckoning or simple odometers, e.g., using wheel encoders, human step detection) during ranging operations.

Nodes can perform a limited set of cryptographic operations using pairwise symmetric keys K, obtained through any of the standard symmetric key establishment protocols (since key establishment has been extensively studied, we make use of existing techniques for key establishment in adhoc networks [23]). Consequently, each pair of nodes, A and B, shares a symmetric key, K_{AB}. Cryptographic operations include encryption, message authentication, and hash computations. Nodes can generate random nonces as needed.

B. Threat Model

The threat consists of a set of static attackers distributed across a geographic region. Each attacker is equipped like a correct node and has similar radio and ranging interfaces. Additionally, each attacker has a second network interface capable of communicating with other attackers using low latency links imperceptible to normal nodes.

Attackers are external [7] and do not have the ability to compromise a correct node. They do not have access to cryptographic keys and are computationally incapable of defeating encryption. Attackers can perform ranging calculations but do not have the ability to know the location of a correct node.

Attackers are organized into wormholes that perform fast-relay attacks [7] in which messages are forwarded at the symbol level. Wormhole activities add negligibly to the total latency of communications. An attacker cannot be continuously either collocated with, or in the immediate vicinity of, a correct node. Otherwise, it would be impossible to distinguish between the location of the mobile node, and the location of the wormhole.

Both RF communications and ranging signals between nodes may be affected by the wormhole. Messages may be selectively delayed or discarded. The wormhole has the ability to modify or replicate communications but encryption and unique nonces limit the impact of these activities.

C. Problem Formulation

Figure 1 provides a framework for understanding the MSND protocol. Node A moves through an area. Node B is also mobile, as shown in Figures 1(a) and 1(b). As nodes come into contact with each other, they attempt to communicate. However, over a wireless connection there is no guarantee that these potential neighbors actually lie within the same neighborhood. Although encryption protects the contents of communications between two nodes, communicating nodes may actually be connected through a wormhole. As shown in Figure 1(c), which resembles the scenario shown in Figure 1(a), a wormhole has the ability to selectively relay, delay or deny communications. Nodes A and B, which are not real neighbors, may be convinced by the wormhole that they are. In order to verify that two communicating nodes are local to the same neighborhood, the nodes conduct MSND.

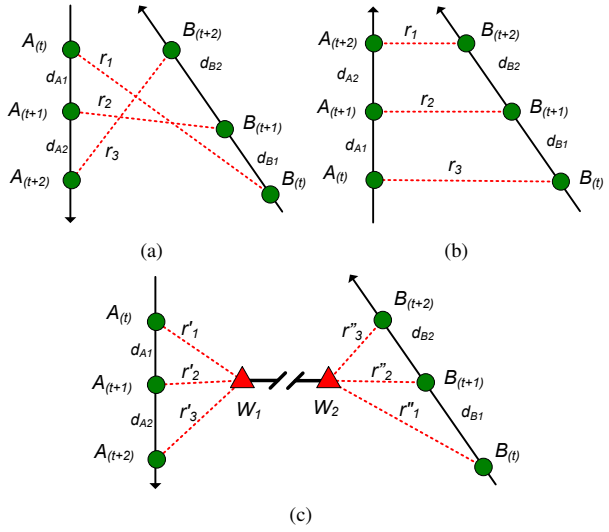

Fig. 1. Neighbor discovery scenarios for two nodes A and B with solid lines showing movement, dashed lines indicating connectivity, and W_1 and W_2 representing the ends of a wormhole. a) Two mobile nodes move in opposite directions. b) Two mobile nodes move in the same direction. c) Two mobile nodes communicate through a wormhole. The range perceived by nodes is $r_i = r'_i + r''_i$.

IV. MSND Protocol

A. Main Idea

The MSND protocol is based on the intuition that when nodes range while moving, the length of the next range is related to the distance traveled between consecutive ranges. Since the wormhole is unable to know the distance traveled by each node, it is not able to influence ranging operations in a way that causes a consistent set of ranges to be built. Graph rigidity is key to this intuition.

In [24], Laman's Theorem states that graph G, composed of rigid edges connected by flexible joints, is minimally rigid in a plane if and only if it has k vertices and $2k-3$ independent edges, and if every induced subgraph on k vertices has at most $2k-3$ edges. When two nodes travel and range, their motion and relative positions can be described as a graph. In many scenarios, the graph produced is rigid and rigidity properties can be leveraged to support wormhole detection and localization.

When two nodes travel along describable paths, it is possible to define one node's line of travel relative to the other. The lines of travel may converge, diverge or be parallel. After two ranges, there is an infinity of possible relationships between the two paths. Three ranges limit the number of relative paths to a few discrete scenarios while four or more produce a rigid graph. In this rigid graph, it is possible to accurately estimate the expected lengths of the next ranges and compare them to the actual ranged value.

In this same movement scenario, a wormhole induces distortion due to its position relative to the lines of travel of each node. When no wormhole is present, a ranging signal travels directly from the sender to the receiver, as shown

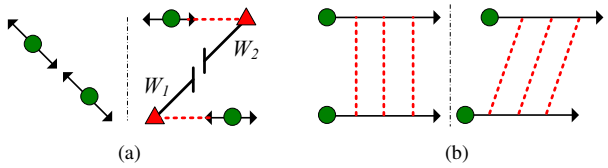

Fig. 2. Degenerate cases. a) All points collinear and all points apparently collinear due to wormhole; b) Lines of travel are parallel and all ranges are equal.

Algorithm 1 MSND Protocol

1: **for** i = 1 to NR **do**
2: $r_i \leftarrow$ range(node A, node B)
3: d_{Ai},$d_{Bi} \leftarrow$ move(node A, node B)
4: **end for**
5: $wh\_present \leftarrow$ Verification
6: **if** false = $wh\_present$ **then**
7: Become neighbors
8: **end if**

in Figure 1(a). However, in the presence of a wormhole, the ranging signal must travel from the sender to the near side of the wormhole, transit the wormhole, and then travel from the distant side of the wormhole to the second node, as shown in Figure 1(c). If ranging nodes were static, this distortion would be impossible to detect with only two nodes. However, mobility causes the distance between each node and its associated end of the wormhole to change. This change in distance ($r_i = r'_i + r''_i$, as shown in Figure 1(c)), translates to ranges that are longer than expected, and to pairs of consecutive ranges whose lengths vary by more than predicted in the rigid graph produced by the nodes' movements.

Of note, however, is that Laman's Theorem applies only to generic frameworks that are not geometrically degenerate [25], and in which the vertices are distributed "wisely" [26]. Demaine [27] notes that most frameworks are generic except for those with very specific and degenerate alignments, like when all points are collinear or when all edges are parallel and of the same length. So, although rigidity is an expected outcome of node movements, there are degenerate cases that impact the MSND protocol. One case is that of two nodes moving along the same line of travel, as noted in Figure 2(a). This produces a graph where all points are collinear. The same result is produced when each node moves directly towards or away from its respective end of the wormhole. Figure 2(b) demonstrates cases where lines of travel are parallel and all ranges are the same length. These graphs are infinitely flexible.

B. MSND Protocol Overview

MSND executes in two phases, as shown in Algorithm 1. First, NR ranging operations are conducted (Lines 1-4) and the resulting ranges and travel distances are sent to Verification (Line 5).

The execution of MSND between two mobile nodes in a network with no wormhole is shown in Figures 1(a) and 1(b). Node A initiates ranging operations. Node B will become a neighbor of A if it can be verified. Figure 1(c) shows a similar scenario in the presence of a wormhole. In the sections that follow, we will use notation of variables consistent with those in Figure 1.

C. Ranging

Ranging consists of three steps similar to [14]. A key requirement of the ranging phase is that each ranging node must travel along a describable path. For the purposes of this paper, nodes move in straight lines until either enough ranges are collected or it is no longer possible to range. Ranging operations stop before completion of the protocol when nodes are no longer in contact or when a node is forced to turn.

The first step, Synchronization, allows participating nodes to calculate the difference in their clocks. The second step, Transmission, provides the ranging signal. The final step, Data Exchange, involves an exchange of data that terminates with both nodes aware of the range between themselves.

Synchronization:

$A \xrightarrow{RF} B : \langle REQ, E_{K_{AB}}\{N_B^r\}, H\{N^S\}, MAC_{K_{AB}}\{.\}\rangle$
$A :$ $t^A_{REQ} :=$ Sending time of REQ
$B :$ $t^B_{REQ} :=$ Reception time of REQ
$B :$ If N_B^r is fresh and MAC is correct then:
$B \xrightarrow{RF} A : \langle REP, N_B^r, , MAC_{K_{AB}}\{.\}\rangle$
$B :$ $t^B_{REP} :=$ Sending time of REP
$A :$ $t^A_{REP} :=$ Reception time of REP

In the Synchronization step of ranging, nodes A and B exchange two packets. Node A sends a request packet containing a nonce encrypted with the pairwise key AB and the hash of a second nonce. The packet is authenticated using a message authentication code generated using pairwise key AB. Node B responds with a packet containing the decrypted nonce that is also authenticated. Both nodes store the transmission and reception time of the two packets.

During the Transmission step, node A ranges by sending a preamble followed by each individual bit of nonce N^s at predetermined intervals. Node B records the arrival time of the preamble and assembles the bits to reconstruct the nonce.

Transmission:

$A \xrightarrow{RNG} * : \langle 1 \parallel N^s\rangle$
$A :$ $t^A_{RNG} :=$ Sending time of RNG
$B :$ $t^B_{RNG} :=$ Reception time of RNG
$B :$ $N^s_B :=$ Received RNG nonce

Data Exchange is the final step. Node A encrypts and sends a packet to node B via RF containing timing information and distance d_A, traveled since the last ranging operation. Nonce N^s is also sent in order to properly associate sets of timing data. Node B stores this data until all ranges are complete and computes its range to A using the ranging signal velocity s (Equation 1).

Data Exchange:

Algorithm 2 Verification

1: $wh\_present?$ ← Conduct preliminary checks
2: **if** $(wh\_present)$ **return** true
3: **for** i = 1 to 3 **do**
4: $\quad \boldsymbol{X} \leftarrow$ get rigid graph $(\boldsymbol{D})$
5: $\quad \tau \leftarrow$ Test fit $(\boldsymbol{X}, y(x))$
6: **end for**
7: $wh\_present \leftarrow$ Vote $(\tau \geq TH$ or $\sigma \geq ST)$
8: **if** $(!wh\_present$ and TestAngle) **then**
9: $\quad$ **if** $(angle(\boldsymbol{X}) \leq AT)$ **return** warning
10: **end if**
11: **return** $wh\_present$

$A:$ If B has sent the correct REP, then:
$A \xrightarrow{RF} B: \langle ACK, E_{K_{AB}}\{N^s, t^A_{REQ}, t^A_{REP}, t^A_{RNG}, d_A\}\rangle$
$B:$ If MAC is correct and $N^s = N^s_B$ and $|(t^A_{REP} - t^A_{REQ}) - (t^B_{REP} - t^B_{REQ})| < \varepsilon$:

$$r_i = ((t^B_{RNG} - t^B_{REQ}) - (t^A_{RNG} - t^A_{REQ})) \times s \qquad (1)$$

D. Verification

Verification uses preliminary checks, metric multidimensional scaling (MDS) and knowledge of node movement to detect distortions caused by a wormhole. Algorithm 2 is used to analyze ranges and traveled distances to determine if a wormhole has affected the results. Successful verification confirms that the two nodes are neighbors.

Verification begins with preliminary checks (Line 1) that include a check for ranges that are too long, adjoining ranges whose length differs by more than the combined distances traveled by the participating nodes, and degenerate configurations. Successful preliminary checks are followed by a loop that performs distance analysis using MDS (Line 4) and a test of the fit of the resulting coordinates (Line 5). The output is analyzed and the best two outcomes are used to make a decision about the presence/absence of a wormhole (Line 7).

The first step of Verification is a set of preliminary checks that analyze the distances for easily detectable evidence of wormhole involvement. Preliminary checks include:

1) $r_i > R_{RNG} + \epsilon$. Ranges as large as $2 \times R_{RNG} + delay$ may be produced by a wormhole. Ranges that exceed R_{RNG} by some defined threshold violate the propagation properties of the ranging signal.
2) $r_{i+1} = R_i \pm d_{Ai} \pm d_{Bi}$. When all points are collinear, the change in length of consecutive ranges is the direct results of adding and/or subtracting node travel distances.
3) $(r_i - d_{Ai} - d_{Bi}) < r_{i+1} < (r_i + d_{Ai} + d_{Bi})$. The length of range r_{i+1} can be no greater than the sum of r_i and the distance traveled by each node between ranges. It can be no smaller than their difference.
4) $(r_i = r_{i+1} = r_{i+2})$ & $(\sum_{i=1}^{r} d_{Ai} = \sum_{i=1}^{r} d_{Bi})$. If all ranges are equal and traveled distances are equal, then the graph produced is not rigid.

Once preliminary checks are complete (as shown in Algorithm 2 Line 4) the ranges r_i and travel distances d_{Ai} and d_{Bi} are passed to MDS in the form of a matrix $\boldsymbol{D}$ of size $2 \times NR$ (as mentioned before, NR is the number of ranges collected). It is important to observe that, for MSND, we do not have distances between any two points in the graph. For example, we do not have a range r_{12} between $A(t)$ and $B(t+1)$ because mobile nodes are not required to stop (and, hence, have $A(t)$ the same as $A(t+1)$). Hence, the problem MDS attempts to solve contains only partial "similarities" between points. More precisely, matrix $\boldsymbol{D}$ is defined as follows:

$$\boldsymbol{D} = \left(\begin{array}{cccc} 0 & A_{ij} & r_i & NaN \\ A_{ij} & 0 & NaN & r_i \\ \hline R_{i-p} & NaN & 0 & B_{i-p,j-p} \\ NaN & r_{i-p} & B_{i-p,j-p} & 0 \end{array}\right)$$

where $A_{ij} = \sum_{k=i}^{j} d_{Ak}$, $B_{ij} = \sum_{k=i}^{j} d_{Bk}$, NaN indicates the absence of a distance between the points. We use p instead of NR, for condensed notation. The steps of classical multidimensional scaling are then [28]:

1) compute the squared distance matrix: $\boldsymbol{D}^{(2)} = [d^2_{ij}]$;
2) double-center the $\boldsymbol{D}^{(2)}$ matrix: $\boldsymbol{B} = -\frac{1}{2}\boldsymbol{J}\boldsymbol{D}^{(2)}\boldsymbol{J}$;
3) compute the singular value decomposition of $\boldsymbol{B} = \boldsymbol{V}\boldsymbol{A}\boldsymbol{V}^T$;
4) compute the coordinate matrix: $\boldsymbol{X} = \boldsymbol{V}_+\boldsymbol{A}_+^{1/2}$, where $\boldsymbol{A}_+$ is the matrix of the first m singular values and $\boldsymbol{V}_+$ the first m columns of $\boldsymbol{V}$.

The output of MDS is $\boldsymbol{X}$, the set of coordinates that describes each node's path of travel. The goodness of MDS output is characterized by a stress factor: $\sigma = \sqrt{\frac{\sum_{ij}(d_{ij} - d^*_{ij})^2}{\sum_{ij} d^2_{ij}}}$. Since mobile nodes have knowledge about their path of travel $(y = f(x))$, MSND will fit y through the node's path of travel, as indicated in Algorithm 2 Line 5. For estimating the goodness of fit we use the norm of residuals: $\tau = \sqrt{\sum_i (y_i - \widehat{y_i})^2}$. Considering ranging and travel errors, the goodness of fit τ is expected to vary. The criteria for determining the existence of a wormhole depends on τ and σ, e.g., if $\tau \geq TH$ or $\sigma \geq ST$, an inconsistency in rigidity of the graph is determined, hence a wormhole attack is signaled. Consequently, different threshold values TH and ST need to be determined, based on the errors expected during mobile node ranging and travel. As a last step (Algorithm 2 Line 8), we test that we are not encountering a degenerate case, i.e., where the trajectories are parallel (as produced by MDS). Hence, if the degenerate case identification is enabled and if we did not detect a wormhole, the algorithm ensures that node trajectories form an angle greater than AT.

V. Security Analysis

In this section we present the security analysis of MSND.

Proposition 1: A wormhole, W_1:W_2, cannot determine the range between two mobile nodes.

Proof: During the ranging portion of MSND, the wormhole forwards a ranging signal from sender to receiver. However, the wormhole does not know the precise time of transmission. The receiver passively receives the signal. Although the MSND protocol requires the exchange of RF packets after ranging, their transmission occurs at an arbitrary interval after ranging signal reception and processing. Without an accurate measure of the signal's time of flight, the wormhole has no ability to determine the range between mobile nodes. ■

Proposition 2: A wormhole, W_1:W_2, cannot determine the distance traveled by a mobile node.

Proof: For the sending node, the only distance information available to the wormhole is metadata related to the ranging signal, i.e., the signal strength. Similarly, the only distance information available to the receiving node is metadata from the RF packet, i.e., the RSSI. However, the velocity of nodes between transmissions is unknown and neither set of metadata reliably produces accurate distance information. ■

Proposition 3: A wormhole, $W_1 : W_2$, cannot determine ranges and distances traveled by reading the contents of the packets it forwards.

Proof: By the System Model, wormholes are unable to break encryption. ■

Theorem 1: MSND is secure.

Proof: Laman's Theorem notes that for a graph $G = (V, E)$ to be generically rigid in the plane, it must have n vertices and $2n - 3$ independent edges. In graphs that have more than $2n - 3$ edges, there must be a subset $F \subseteq E$ that satisfies two conditions: (1) $|F| = 2n - 3$ and (2) for all $F' \subseteq F, F' \neq \emptyset, |F'| \leq 2k - 3$, where k is the number of vertices which are endpoints of edges in F' [26].

In MSND, the number of edges $E = r^2$ where r is the number of ranges (the edges of the graph are the ranges, distances traveled and combinations of distances traveled). So, after the third range, $E > 2n - 3$ where $n = 2r$. However, in every subgraph F where $r = 3$, then $E = 2k - 3$ and for every $F' \subseteq F, |F'| \leq 2k - 3$. So, by Laman's Theorem, the graph produced by node movements is generically rigid.

In a rigid graph, the length of the range, can be predicted if previous range and distances traveled are known. As noted in the preliminary checks, a wormhole cannot shorten a range. However, by Proposition 1, the wormhole cannot know the value of the previous range. Delaying the signal arbitrarily might violate a preliminary check. Therefore, the wormhole must select a delay that will produce ranges that embed in a rigid graph. However, by Proposition 2, the wormhole does not know the distance traveled by a node between ranging signals. The only source of this information is encrypted which, by Proposition 3, is not available to the wormhole.

Since the wormhole cannot know the distances traveled between ranges by nodes, the lengths $r^2 - r$ edges are unknown. The lengths of the remaining r edges are also unknown because the wormhole cannot determine a range. Therefore, any non-degenerate graph affected by the wormhole will not be rigid. ■

VI. Performance Evaluation

Our simulation experiments were conducted using a purpose-built simulator. Movement and ranging were handled in a Java simulation of a random waypoint movement model. Nodes move in a 900×300 field populated by a single wormhole with two ends. Both R_{RF} and R_{RNG} are set to 300. At each waypoint, node speed is chosen randomly between 2%-7% of R_{RNG}. Ranges and travel distances, in the form of matrix $\boldsymbol{D}$, were used in a Matlab implementation of metric MDS. The output of each of three MDS iterations was processed through the $polyfit$ function and a "best two of three" voting algorithm made the wormhole attack decision.

Each experimental point is the average of 10 simulation runs with 10 MSND verifications per run, for a total of 100 verifications. Each run used a different pseudo-random seed. Key metrics for evaluating MSND are false positives (FP) and false negatives (FN) in wormhole detection. False positives are scenarios in which MSND indicated a wormhole was present when there was none. False negatives are scenarios in which MSND fails to indicate that a wormhole is present, when there was one. False negative results are the most important. While a false positive outcome is erroneous, security is not compromised. However, a false negative outcome means that nodes that are not real neighbors become neighbors through a wormhole.

We investigate the effect of the following parameters: number of ranges r_i used (denoted by NR), ranging error in r_i (denoted by RE), threshold τ selection (denoted by TH). We used a stress threshold $ST = 0.001$.

A. Number of Ranges

For this experiment, RE was set to zero while $TH = 2$. As evidenced by the results, presented in Figure 3(a), system performance is robust to the different number of ranges. We observe that FP and FN do not improve significantly with an increased number of ranges, suggesting that waiting for additional ranging operations to complete would not benefit the accuracy of MSND. For our simulation, when NR=12, the results are FN=10% and FP=15%.

B. Threshold Selection

The criteria to verify neighbors is a threshold value TH for the norm of residuals τ, as output from a linear fitting function to the set of coordinates produced by MDS. The selection of TH, above which a wormhole is reported, is an important and tunable parameter that might vary for several reasons. As shown in Figure 3(b), FN and FP are inversely proportional. A change in TH that lowers FP will raise FN. Selecting a very low TH dramatically reduces FN but has the additional consequence of maximizing FP. Selection of TH for a specific network depends on the degree of error in ranging and travel measurements, tolerance for FN, and an assessment of how often nodes must communicate.

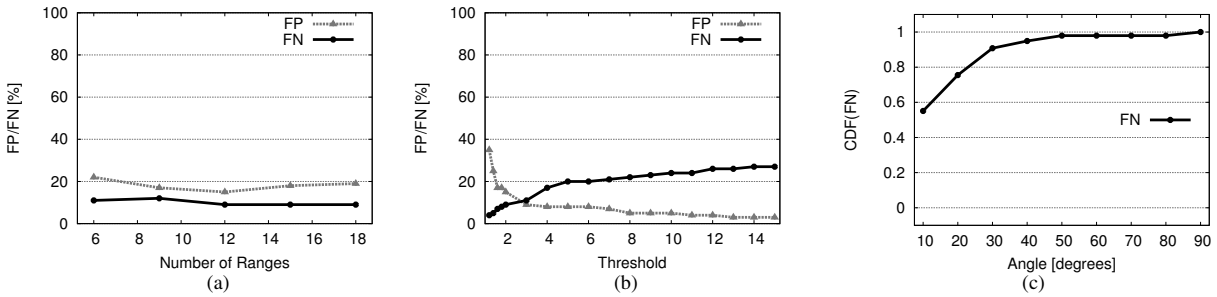

Fig. 3. (a) Effects of number of ranges on system performance (RE=0.0, TH=2); (b) Effects of threshold on system performance (RE=0.0, NR=12); (c) Cumulative distribution for FN as a function of the angle between mobile node trajectories, as obtained by MDS.

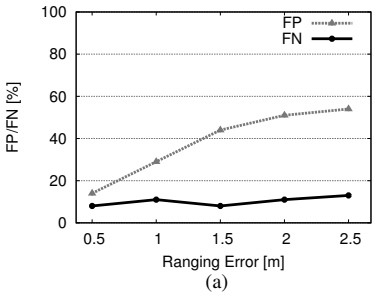

Fig. 4. Effects of range error on system performance (NR=12 TH=[4,5,9,9,13]).

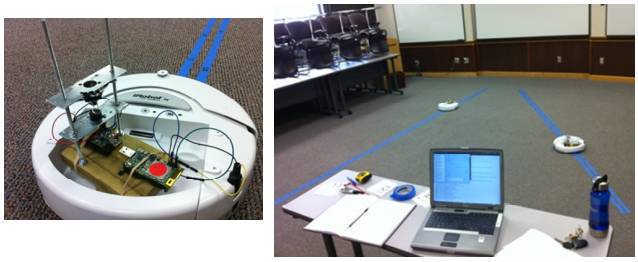

Fig. 5. Experimental evaluation using Epic motes ranging with SRF-02 ultrasonic range finders and carried by iRobots. A mobile node equipped with ranging hardware, and experimental setup.

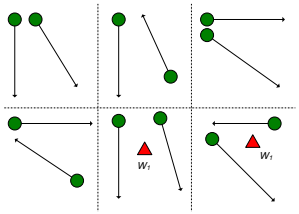

Fig. 6. The six scenarios executed in the experiment. Wormholes are represented with red triangles.

C. Angle Analysis

In this section we investigate the effect the angle between the mobile node trajectories (as obtained by MDS) in $\boldsymbol{X}$ (Algorithm 2 line 8) has on performance of MSND, and justify the check for its value to be less than a threshold AT. In Figure 3(c) we present the CDF for false negatives FN as a function of the angle produced by MDS. Remarkably, most of the false negatives (i.e., over 50% of total FN) are when MDS predicts that the trajectories are close to parallel. The degenerate scenarios that MDS sometimes identifies become problematic especially when ranging errors are encountered. We evaluate its effects on MSND in the following sections.

D. Ranging Error

The impact of ranging error RE is shown in Figure 4(a). RE is calculated in absolute units and it is assumed to have a Gaussian distribution with variance as indicated in the figure. TH was increased with increasing error with specific TH values noted in the caption. As error increases, the total number of FN remains roughly constant. False-positives, however, increase sharply. Analysis of FP outcomes revealed that many of the failures results were from scenarios where one of the nodes selected a very slow minimum speed and, thus, traveled a short total distance.

VII. System Implementation

We implemented MSND on EPIC motes running TinyOS 2.1.1 which were hosted on iRobot Create robots, as shown in Figure 5. Ranging was conducted using Devantech SRF-02 ultrasonic rangefinders interfaced with the motes via UART. Verification was handled by a centralized node. The wormhole was emulated by placing a third Epic mote equipped with an ultrasonic rangefinder at a point located between the two nodes' lines of travel, as shown in Figure 6. Ranging was conducted independently between the wormhole and each node (ranges r_i' and r_i''). The resulting two values were combined to determine the total length of the range between the two nodes via the wormhole, i.e., $r_i = r_i' + r_i''$.

We performed experiments in a $10\times10\text{m}^2$ indoor office environment. As shown in Figure 6, six scenarios involving two nodes were executed, out of which two involved a static wormhole. Each node traveled an average of 5m per scenario in a straight line. The triangles in Figure 6 denote the approximate position of the wormhole in relation to the nodes' line of travel.

For each scenario, the iRobots took 12 steps of around 30-40cm each. Some errors were introduced by the robot's distance measuring mechanism, resulting in an actual step size

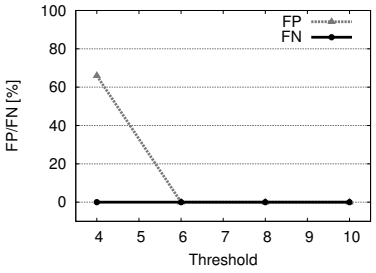

Fig. 7. Results of the experiment. Note that FN is always 0.

of $\pm 3\%$ of the original. At each step, the nodes ranged with an accuracy of $\pm 1.49\%$. The actual distance between nodes and between steps was recorded manually, for ground truth.

A. System Evaluation

Figure 7 shows that the wormhole was always detected, even at high TH. Given the independent travel and ranging errors noted in the system implementation section, minimum predicted TH values lie in the range of 6-9 for ranging error and 5-6 for travel error. Remarkably, the FN rate was 0%. In fact, in the two wormhole scenarios, the minimum TH that would allow an FN is 32.5, a comfortable margin above TH=18, the point where all no-wormhole cases were identified.

VIII. Conclusions and Future Work

The ability to securely determine valid neighbors is an important part of many network functions. In a network with wormholes, failure to protect neighbor discovery could lead to information disclosure, incorrect localization, routing problems, and adversary control of the network at any time. Increased exposure to DoS attacks may also result. To the best of our knowledge, this paper is the first one to propose a protocol for detecting the presence of a wormhole in scenarios both nodes are mobile. MSND leverages graph rigidity to aid in the verification of network neighbors. An accompanying security analysis demonstrates the secureness of the protocol against a variety of attacks that could be launched by the wormhole including attacks that delay/discard/modify packets. Ongoing and future work will include algorithm enhancements for improved false negative and false positive rates and combination of this work with a localization protocol.

Acknowledgements: This work was funded in part by NSF grant CNS 0923203. The authors thank Stephen George for help with some implementation and editing of the paper, and Sejun Song for feedback.

References

[1] J. Hwang, T. He, and Y. Kim, "Secure localization with phantom node detection," *Ad Hoc Networks*, vol. 6, no. 7, pp. 1031 – 1050, 2008.

[2] P. Papadimitratos and Z. Haas, "Secure routing for mobile ad hoc networks," in *Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS)*, 2002.

[3] S. Vasudevan, J. Kurose, and D. Towsley, "Design and analysis of a leader election algorithm for mobile ad hoc networks," *IEEE Conference on Network Protocols (ICNP)*, 2004.

[4] J. Liu, D. Sacchetti, F. Sailhan, and V. Issarny, "Group management for mobile ad hoc networks: design, implementation and experiment," in *International Conference on Mobile Data Management (MDM)*, 2005.

[5] T. van Dam and K. Langendoen, "An adaptive energy-efficient MAC protocol for wireless sensor networks," in *ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2003.

[6] Y. Hu, A. Perrig, and D. Johnson, "Packet leashes: a defense against wormhole attacks in wireless networks," in *International Conference on Computer Communications (Infocom)*, 2003.

[7] P. P. Papadimitratos, M. Poturalski, P. Schaller, P. Lafourcade, D. Basin, S. Capkun, and J.-P. Hubaux, "Secure Neighborhood Discovery: A Fundamental Element for Mobile Ad Hoc Networking," *IEEE Communications Magazine*, vol. 46, no. 2, 2008.

[8] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location-support system," in *Conference on Mobile Computing and Networking (Mobicom)*, 2000.

[9] H. Chen, W. Lou, and Z. Wang, "A consistency-based secure localization scheme against wormhole attacks in WSNs," in *Conference on Wireless Algorithms, Systems, and Applications (WASA)*, 2009.

[10] R. Poovendran and L. Lazos, "A graph theoretic framework for preventing the wormhole attack in wireless ad hoc networks," *Wirel. Netw.*, vol. 13, no. 1, pp. 27–59, 2007.

[11] S. Capkun, L. Buttyan, and J. Hubaux, "SECTOR: secure tracking of node encounters in multi-hop wireless networks," in *ACM Workshop on Security of Ad Hoc and Sensor Networks*, 2003.

[12] M. Poturalski, P. Papadimitratos, and J. Hubaux, "Secure neighbor discovery in wireless networks: formal investigation of possibility," in *ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2008.

[13] ——, "Towards provable secure neighbor discovery in wireless networks," in *ACM Workshop on Formal Methods in Security Engineering (FMSE)*, 2008.

[14] R. Shokri, M. Poturalski, G. Ravot, P. Papadimitratos, and J. Hubaux, "A practical secure neighbor verification protocol for wireless sensor networks," in *ACM Conference on Wireless Network Security (WiSec)*, 2009.

[15] D. Moore, J. Leonard, D. Rus, and S. Teller, "Robust distributed network localization with noisy range measurements," in *ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2004.

[16] Y. Xu, Y. Ouyang, Z. Le, J. Ford, and F. Makedon, "Analysis of range-free anchor-free localization in a wsn under wormhole attack," in *ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, 2007.

[17] R. Maheshwari, J. Gao, and S. Das, "Detecting wormhole attacks in wireless networks using connectivity information," in *International Conference on Computer Communications (Infocom)*, 2007.

[18] S. Nawaz and S. Jha, "A graph drawing approach to sensor network localization," in *IEEE Conference on Mobile Adhoc and Sensor Systems (MASS)*, 2007.

[19] L. Buttyan, L. Dora, and I. Vajda, "Statistical wormhole detection in sensor networks," in *Security and Privacy in Ad-hoc and Sensor Networks*, 2005.

[20] D. Liu, "Protecting neighbor discovery against node compromises in sensor networks," in *Conference on Distributed Computing Systems (ICDCS)*, 2009.

[21] L. Hu and D. Evans, "Using directional antennas to prevent wormhole attacks," in *Conference on Mobile Computing and Networking (Mobicom)*, 2004.

[22] S. M. George, W. Zhou, H. Chenji, M. Won, Y. Lee, A. Pazarloglou, R. Stoleru, and P.Barooah, "DistressNet: a wireless adhoc and sensor network architecture for situation management in disaster response," *IEEE Communications*, vol. 48, no. 3, Mar. 2010.

[23] Z. Li and J. J. Garcia-Luna-Aceves, "Non-interactive key establishment in mobile ad hoc networks," *Ad Hoc Netw.*, vol. 5, 2007.

[24] G. Laman, "On graphs and rigidity of plane skeletal structures," *Journal of Engineering Mathematics*, vol. 4, no. 4, pp. 331–340, Oct. 1970.

[25] I. Streinu and L. Theran, "Combinatorial genericity and minimal rigidity," in *Symposium on Computational Geometry (SCG)*, 2008.

[26] B. Servatius and H. Servatius, "Generic and abstract rigidity (private communication)," 1999.

[27] E. D. Demaine and J. O'Rourke, *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, 2008.

[28] I. Borg and P. Groenen, *Modern Multidimensional Scaling. Series in Statistics.* Springer, 1997.