Statistical Wormhole Detection in Sensor Networks

Levente Buttyán, László Dóra, and István Vajda

Laboratory of Cryptography and System Security (CrySyS),
Department of Telecommunications,
Budapest University of Technology and Economics, Hungary
{buttyan, laszlo.dora, vajda}@crysys.hu

Abstract. In this paper, we propose two mechanisms for wormhole detection in wireless sensor networks. The proposed mechanisms are based on hypothesis testing and they provide probabilistic results. The first mechanism, called the Neighbor Number Test (NNT), detects the increase in the number of the neighbors of the sensors, which is due to the new links created by the wormhole in the network. The second mechanism, called the All Distances Test (ADT), detects the decrease of the lengths of the shortest paths between all pairs of sensors, which is due to the shortcut links created by the wormhole in the network. Both mechanisms assume that the sensors send their neighbor list to the base station, and it is the base station that runs the algorithms on the network graph that is reconstructed from the received neighborhood information. We describe these mechanisms and investigate their performance by means of simulation.

1 Introduction

Sensor networks [1] consist of a large number of sensors that monitor the environment, and a few base stations that collect the sensor readings. The sensors are usually battery powered and limited in computing and communication resources, while the base stations are considered to be more powerful. In order to reduce the overall energy consumption of the sensors, it is conceived that the sensors send their readings to the base station via multiple wireless hops. Hence, in a sensor network, the sensor nodes are responsible not only for the monitoring of the environment, but also for forwarding data packets towards the base station on behalf of other sensors.

In order to implement the above described operating principle, the sensors need to be aware of their neighbors, and they must also be able to find routes to the base station. An adversary may take advantage of this, and may try to control the routes and to monitor the data packets that are sent along these routes [6]. One way to achieve this is to set up a *wormhole* in the network. A wormhole is a dedicated connection between two physical locations. The adversary installs a radio transceiver at each end of the connection, and it sends and re-transmits every data packet received at one end of the wormhole at the

R. Molva, G. Tsudik, and D. Westhoff (Eds.): ESAS 2005, LNCS 3813, pp. 128–141, 2005.

other end of it. Practically, this means that the adversary creates communication links between some pairs of sensors that would otherwise not be able to communicate directly with each other. In other words, the adversary modifies the topology of the network. If this is done carefully, the adversary may achieve that many sensors send their data packets to the base station via the wormhole. While the application of cryptographic mechanisms (e.g., encryption and message authentication) prevents an adversary from monitoring and modifying the information sent to the base station, cryptographic mechanisms do not solve every problem stemming from wormholes. The adversary can still mount denial of service type attacks, such as dropping packets (possibly selectively) that are transferred through the wormhole. In addition, the sensors which are close to the transceivers of the wormhole participate more in packet forwarding and they deplete their battery earlier. Therefore, in most of the applications, wormhole detection is an important requirement.

In this paper, we propose two mechanisms for wormhole detection in wireless sensor networks. The proposed mechanisms are based on hypothesis testing and they provide probabilistic results. The first mechanism, called the Neighbor Number Test (NNT), detects the increase in the number of the neighbors of the sensors, which is due to the new links created by the wormhole in the network. The second mechanism, called the All Distances Test (ADT), detects the decrease of the lengths of the shortest paths between all pairs of sensors, which is due to the shortcut links created by the wormhole in the network. Both mechanisms assume that the sensors send their neighbor list to the base station, and it is the base station that runs the algorithms on the network graph that is reconstructed from the received neighborhood information. The main advantage of the proposed mechanisms is that they do not require special hardware in the sensors, directional antennas, tight clock synchronization, or distance measurements between the nodes. The only requirement is that the sensor nodes can determine who their neighbors are, and they can send this information to the base station in a secure way. The rest of the paper is organized as follows. In Section 2, we overview the state-of-the-art in the filed of wormhole detection. In Section 3, we present our approach by describing the operation of the two wormhole detection mechanisms that we propose. The effectiveness of the mechanisms is studied in Section 4, where we present and analyze our simulation results. Finally, in Section 5, we conclude the paper and sketch some possible future research directions.

2 Related Work

In [5], the authors propose two approaches for detecting wormholes in wireless ad hoc networks, where sensors are allowed to move during the communication. The first approach is called *geographical packet leashes*, and it requires the nodes to be aware of their own location and to maintain loosely synchronized clocks. Every time when a node A sends a packet to its neighbor B, it puts its location and the time of sending into the header of the packet. When the packet is received by

B, it compares the time of reception to the time of sending, and calculates the maximum distance between A and B using the difference between their locations and the distance that they could move away between sending and receiving the packet. If the estimated distance is longer than the possible maximum radio range then B rejects the communication with A.

The other approach is called *temporal packet leashes*, and it avoids using any special hardware for localization, but it requires tightly synchronized clocks. Every time when a node A sends a packet to its neighbor B, it puts an authenticated time stamp into the header. When B receives the packet, it calculates the possible maximum distance between A and B from the difference between the time of sending and the time of receiving of the packet, and assuming that the packet travels with the speed of light. If the resulting distance is too large, then this indicates a wormhole. This procedure relies on the fact that going through the wormhole means covering a longer distance than the normal distance between neighboring nodes, and this longer distance can be precisely measured due to the tightly synchronized clocks.

The disadvantage of the above approaches is that they require either location information of each node or tight clock synchronization between the nodes, and these requirements cannot always be satisfied in sensor networks.

In [3], another approach is proposed to estimate the real physical distance between two communicating nodes, which does not require location information or clock synchronization at all. That approach is based on an authenticated *distance-bounding* protocol, called MAD. The distance-bounding phase of MAD consist of several rounds, and in each round, each node sends a one bit challenge to the other node to which the other node responds with a one bit response immediately. Each node locally measures the time between sending out the challenge and receiving the response, and based on the measured times, it estimates its distance to the other node, assuming that messages travel with the speed of light. In order for this to work, the nodes must be able to measure local timings with nanosecond precision, which is possible with today's hardware. In addition, it is crucial that the response is sent immediately after receiving the challenge. This, however, may not be possible using standard hardware. The main problem is that typical wireless medium access control protocols introduce random delays between the time at which the application sends a message and the time at which that message is really transmitted via the radio interface. Therefore, this approach also requires special hardware in the sensor nodes and special medium access control protocols.

Another wormhole detection approach that uses the node's location information is proposed in [7]. However, as opposed to the geographical leash approach proposed in [5], here only a small fraction of the nodes need to be equipped with a GPS receiver. These special nodes are called *guards* and it is also assumed that the guards have a larger radio range (denoted by R) than the other nodes. The guards broadcast their positions in their one hop neighborhood. Two nodes consider each other neighbor only if they hear a threshold number of common guards. The nodes use the location information broadcast by the guards

to detect wormholes based on the following two principles: (i) since any guard heard by a node must lie within a range of radius R around the node, a node cannot hear two guards that are $2R$ apart from each other; and (ii) since the messages sent by the guards are authenticated and protected against replay, a node cannot receive the same message twice from the same guard. It is shown in [7] that based on these principles, wormholes can be detected with probability close to one. However, the disadvantage of this approach is that the guards are distinguished nodes in the network that differ from the regular nodes.

In [4], the authors propose a wormhole detection approach that assumes that the nodes know from which direction they got a packet. The intuitive idea behind this approach is that if there is no wormhole in the system, then the following must be true: if one node sends a packet in a given direction, then its neighbor will hear that packet from the opposite direction. However, if there is a wormhole in the system, then the above statement is not always true (depending on the placement of the wormhole), and thus, the wormhole becomes detectable. Unfortunately, it has a significant probability that the wormhole is there, but it is not caught. In order to address this problem, the authors worked out two algorithms in which the nodes involve their neighbors during the communication to help to discover the wormhole. The main disadvantage of this approach is that it requires directional antennas, which are usually not available in sensor networks.

In [8], a centralized wormhole detection technique is proposed, which uses inaccurate distance estimations between neighboring nodes. The main idea of the proposed technique is to reconstruct a virtual layout of the network and identify inconsistencies in it. For this reason, the connectivity information and the inaccurately estimated distances between the neighbors are fed into a multidimensional scaling (MDS) algorithm, which tries to determine a virtual position for every node in such a way that the constraints induced by the connectivity and the distance estimation data are respected. Since the distances are estimated inaccurately, the algorithm has a certain level of freedom in "stretching" the nodes within the error bounds of the distance estimation. If the estimated distance between two nodes connected by a wormhole are much larger than the nodes' communication range, then the wormhole is detected immediately. Hence, the adversary must falsify the distance estimation and arrange that the estimated distances between the nodes affected by the wormhole become credible. However, this will result in a distortion in the virtual layout constructed by the MDS algorithm; in particular, the layout will be contracted between the affected nodes. By visualizing the virtual layout or by computing appropriate indicator values, the distortion can be detected and the wormhole can be located.

3 Our Approach

Compared to the above described approaches, our approach neither requires special hardware and directional antennas in the nodes, nor tight clock synchronization and distance measurements. We only assume that the sensor nodes can

determine who their neighbors are, and they can send this information to the base station(s). Based on the received neighborhood information, the base station(s) can detect the presence of wormholes probabilistically using hypothesis testing. In this section, we propose two specific mechanisms for this purpose; we will evaluate the effectiveness of the proposed mechanisms in Section 4.

3.1 System Assumptions

We assume that the system consists of a large number of sensor nodes and a few base stations placed on a two dimensional surface. We assume that the base stations have no resource limitations, and they can run complex algorithms. We assume that the sensors have a fixed radio range r, and two sensors are neighbors, if they reside in the radio range of each other. We assume that the sensors run some neighbor discovery protocol, and they can determine who their neighbors are. We also assume that the sensors send their neighborhood information to the closest base station regularly in a secure way. By security we mean confidentiality, integrity, and authenticity; in other words, we assume that the adversary cannot observe and change the neighborhood information sent to the base stations by the sensors, neither can it spoof sensors and fabricate false neighborhood updates. This can be ensured by using cryptographic techniques that we will not detail in this paper. Note that the neighborhood information can be piggy-backed on regular data packets. In addition, as sensor networks tend to be rather static, sending only the changes in the neighborhood since the last update would reduce the overhead significantly. The base stations can pool the received neighborhood information together, and based on that, they can reconstruct the graph of the sensor network. We assume that the node density is high enough so that the network is always connected.

We assume that the adversary can set up a wormhole in the system. The wormhole is a dedicated connection between two physical locations. There are radio transceivers installed at both ends of the wormhole, and packets that are received at one end can be sent to and re-transmitted at the other end. In this way, the adversary can achieve that nodes that otherwise do not reside in each other's radio range can still hear each other and establish a neighbor relationship (i.e., they can run the neighbor discovery protocol). This means that the adversary can introduce new, otherwise non-existing links in the network graph that is constructed by the base stations based on the received neighborhood information.

The wormhole is characterized by the distance between the two locations that it connects and the radio ranges of its transceivers. We assume that the receiving and the sending ranges of both transceivers are the same, and we will call this range the *radius* of the wormhole. The radius of the wormhole is not necessarily equal to the radio range of the sensors.

In principle, the adversary can drop packets carrying neighborhood information that are sent to the base stations via the wormhole. However, consistently missing neighborhood updates can be detected by the base stations and they indicate that the system is under attack. Therefore, we assume that the adversary does not drop the neighborhood updates. In addition, by the assumptions made earlier, it cannot alter or fabricate them either.

3.2 Neighbor Number Test (NNT)

Our first detection mechanism is based on the fact that by introducing new links into the network graph, the adversary increases the number of neighbors of the nodes within its radius. This is illustrated in Figure 1. The thick circle in the figure is the radio range of the sensor node A. Its real neighbors are N_i within the radio range of sensor node A. The two other circles show the radius of the wormhole. The nodes at the further end of the wormhole that are labelled with W_i are the neighbors that are due to the existence of the wormhole. These sensors are outside of the radio range of A, and they would not be its neighbors if there was no wormhole.

If the distribution of the placement of the nodes is given, then it is possible to compute the hypothetical distribution of the number of neighbors. Then, the base stations can use statistical tests to decide if the network graph constructed from the neighborhood information that is received from the sensors corresponds to

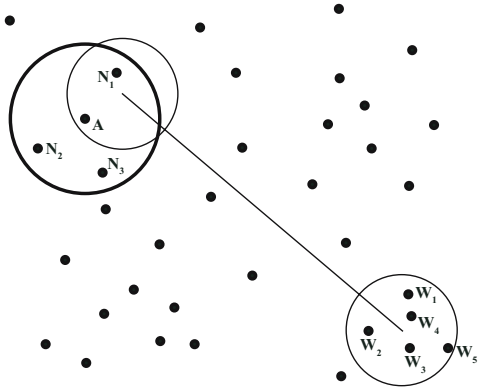

Fig. 1. The wormhole increases the number of neighbors of the nodes in its radius

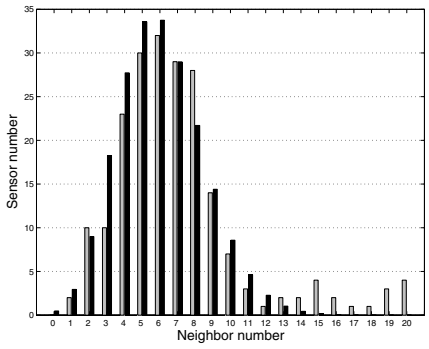

Fig. 2. Hypothetical (dark) and real (light) distributions of the number of neighbors

this hypothetical distribution. In order to illustrate this idea, let us consider the example depicted in Figure 2, where the dark bars correspond to the hypothetical distribution of the number of neighbors, and the light bars show the actual distribution in the network graph reconstructed from the sensors' neighborhood updates. One can see that the probability of higher neighbor numbers (15-20) is increased with respect to the hypothetical distribution, and the idea of the proposed mechanism is to detect this increase by using statistical tests.

Based on the above observations, the NNT algorithm is given as follows:

1. The base station computes the expected histogram of the neighbor numbers using the hypothetical distribution of the number of neighbors.
2. The base station collects the neighborhood updates from the sensors, constructs the network graph, and computes the histogram of the real neighbor numbers in the graph.
3. The base station compares the two histograms with the χ^2–test.
4. If the computed χ^2 number is larger than a preset threshold that corresponds to a given significance level, then a wormhole is indicated.

Computing the parameters for the χ^2–test. Assuming that the sensors are placed uniformly at random on the plane, the probability of two nodes being neighbors is

$$q = \frac{r^2 \cdot \pi}{T}$$

where r is the radio range of the sensor nodes and T is the sphere of the area where the sensor network is deployed. The probability $p(k)$ of having exactly k neighbors is

$$p(k) = \binom{N}{k} \cdot q^k \cdot (1-q)^{N-k}$$

where $N+1$ is the total number of nodes in the network. Let us partition the set $\{0, 1, 2, \ldots\}$ into subsets $B_1, B_2, \cdots, B_m$, such that $e(i) = (N+1)\sum_{k \in B_i} p(k)$ be larger than 5 (a requirement needed by the χ^2–test [2]). The χ^2 number is then computed using the following formula:

$$\chi^2 = \sum_{\forall i} \frac{r(i) - e(i)}{e(i)}$$

where $r(i)$ is the real number of nodes with number of neighbors in B_i. If χ^2 is below the threshold that corresponds to a given significance level (this threshold can be looked up in published tables of χ^2 values), then the hypothesis is accepted, and no wormhole is indicated. Otherwise the hypothesis is rejected, and a wormhole is indicated.

3.3 All Distances Test (ADT)

Our second detection mechanism is based on the fact that the wormhole shortens the paths in the network, or more precisely, it distorts the distribution of the

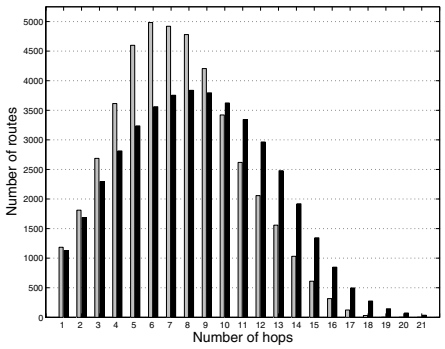

Fig. 3. Hypothetical (dark) and real (light) distributions of the length of the shortest paths between all pairs of nodes

length of the shortest paths between all pairs of nodes. This is illustrated by the example depicted in Figure 3, where the dark bars represent the hypothetical distribution of the length of the shortest paths and the light bars represent the real distribution. As it can be seen, the two distributions are different, and in the real distribution, shorter paths are more likely than in the hypothetical one. The idea is to detect this difference with statistical tests.

The ADT algorithm is very similar to the NNT algorithm:

1. The base station computes the histogram of the length of the shortest paths between all pairs of nodes in the hypothetical case when there is no wormhole in the system using the knowledge of the distribution of the node placement.
2. The base station collects the neighborhood information from the sensors, and computes the histogram of the length of the shortest paths in the real network.
3. The base station compares the two histograms with the χ^2–test.
4. If the computed χ^2 number is larger than a preset threshold that corresponds to a given significance level, then a wormhole is indicated.

Computing the parameters for the χ^2–test. In this case, we were not able to derive a close formula that describes the hypothetical distribution of the length of the shortest paths. Instead, we propose to estimate that distribution by randomly placing nodes on the plane according to the distribution of the node placement, and compute the lengths of the shortest paths between all pairs of nodes in the resulting graph. We propose to repeat the experience many times and average the normalized histograms obtained in these experiences. Once the hypothetical distribution is estimated in this way, the χ^2–test can be used in a similar way as we described in Subsection 3.2.

4 Simulation Environment

In order to evaluate the effectiveness of the proposed mechanisms, we built a simulator that places 300 sensor nodes uniformly at random on a 500 m $\times$ 500 m

Table 1. Simulation parameters

Number of nodes	300
Extent of territory	500 m × 500 m
Number of simulation runs	100
Radio range of sensor nodes	40 m, 47 m, 54 m, 60 m, 65 m, 70 m
Radio range of the wormhole	16 m, 50 m
Distance between the affected areas at the two end wormhole	20 m, 50 m, 100 m, 200 m, 300 m, 400 m

flat area with one base station in the middle, and it also places a wormhole randomly in the same area. The simulator permits us to set three parameters: the radio range of the sensors, the radius of the wormhole, and the distance between the affected areas at the two ends of the wormhole.

We chose two extreme values for the radio range of the sensor nodes: 40 m and 70 m. The expected neighbor number is 5.9 in the 40 m case, and 18.5 in the 70 m case. Then, we split up the range between 5.9 and 18.5 evenly into 5 pieces to get the six radio range values that we used in our simulations (see Table 1).

We set the radius of the wormhole to 16 m or to 50 m (see Table 1). These two values have been selected in such a way that the number of nodes affected by the wormhole differs significantly in the two cases. When the radius of the wormhole is 16 m, one node is affected (falls in the wormhole's range) on both ends of the wormhole on average, whereas when the radius of the wormhole is 50 m, 9.4 nodes are affected on both ends on average.

Finally, we varied the distance between the affected areas at the two ends of the wormhole between 20 m and 400 m (see Table 1).

A given combination of the possible parameter values define a test case. For each test case we run 100 simulations and averaged the results. For each radio range setting, we first determined the rate of the false positive alarms (i.e., the percentage of the simulation runs where the algorithms indicate a wormhole when there is no wormhole in the system). Then, we placed wormholes with different parameters in the system and determined the accuracy of both of our detection mechanisms (i.e., the percentage of simulation runs where the wormhole is detected when there is indeed a wormhole in the system). The results are presented in the following subsections.

4.1 Results of the Neighbor Number Test (NNT)

The results of the NNT algorithm are shown on Figures 4 and 5. Figure 4(a) shows the accuracy of the detection as a function of the radio range of the sensors when the radius of the wormhole is 50 m. As it can be seen, the detection accuracy decreases as the sensors' radio range increases. The reason is that in the case of larger radio ranges, the sensors have more real neighbors, and therefore, the increase in the number of neighbors caused by the wormhole becomes less significant, and consequently, more difficult to detect. We can also observe that the detection accuracy is better when the areas affected by the wormhole are

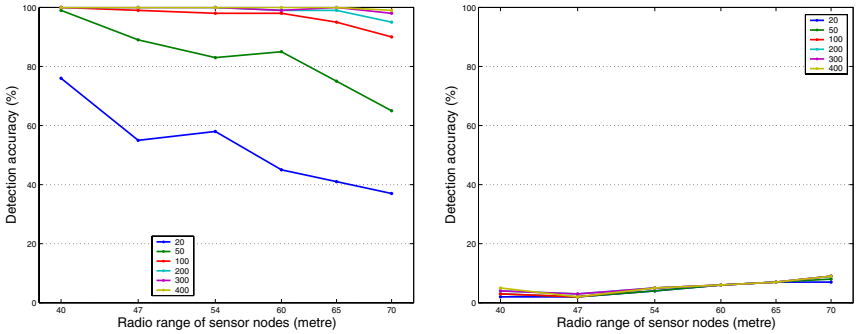

Fig. 4. Detection accuracy plotted against the radio range of the sensor nodes. The different curves belong to different distances between the areas affected by the wormhole with a radius of 50 m (a) and 16 m (b).

more distant from each other, although increasing this distance above 100 m has no real influence on the results. In fact, if the distance between the affected areas is smaller than the radio range of the sensors, then it is possible that two affected nodes that do not belong to the same affected areas are already real neighbors, and therefore, the wormhole does not create a new link between them. In other words, the larger the distance between the affected areas is, the higher the probability is that the wormhole introduces new links into the graph, and by doing so it increases the number of neighbors of the affected nodes.

Figure 4(b) shows the accuracy of the detection as a function of the radio range of the sensors when the radius of the wormhole is 16 m. It is clear from the figure that the NNT algorithm does not work in this case, as the accuracy of the detection is unacceptably low. The huge difference between the performance in the 50 m case and that in the 16 m case can be explained with the large difference in the number of the affected nodes in the two cases. As we described earlier, when the radius of the wormhole is 16 m, on average one node is affected at both ends on the wormhole. Hence, practically, such a wormhole creates a single new link in the graph, which is extremely difficult detect with statistical techniques. On the other hand, as the average number of affected nodes is around 10 at both ends of the wormhole when the radius is 50 m, the number of new links introduced in the graph is around 100. More importantly, around 20 nodes out of the total of 300 have around 10 more neighbors due to the wormhole, and this can be detected by the NNT algorithm.

Figure 5 shows the percentage of the false positive alarms as a function of the radio range of the sensors. As it can be seen, the NNT algorithm performs quite well regarding the false positive alarms. Indeed, the percentage of the false positive alarms is determined by the selected significance level of the χ^2–test, which in our case was 0.025.

In summary, the NNT algorithm detects the wormhole reasonably well if the radius of the wormhole is comparable to or larger than the radio range of the sensors, but it performs very badly if the radius of the wormhole is small. We

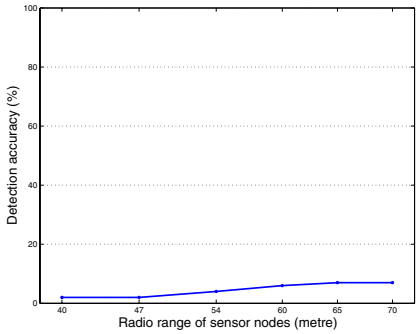

Fig. 5. Percentage of false positive wormhole detections plotted against the radio range of sensor nodes

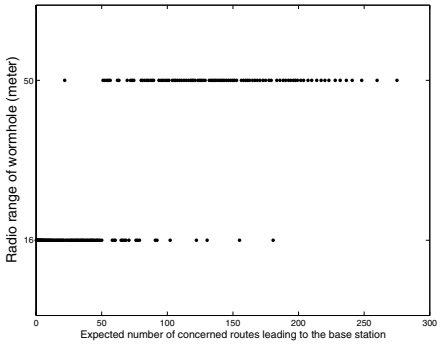

Fig. 6. The effect of the wormhole on the number of the controlled shortest paths plotted against the radius of the wormhole

note, however, that a smaller wormhole radius has smaller effect on the system in terms of the number of sensors that send measurement data to the base station through the wormhole. In order to illustrate this, we constructed the minimum spanning tree rooted at the base station, and counted the number of shortest paths between the base station and the sensors that contain a link created by the wormhole. The result is shown in Figure 6. As it can be seen, when the radius of the wormhole is 16 m, the number of concerned paths is between 0 and 50, whereas in the case of a 50 m radius, the number of concerned paths is between 100 and 200. Thus, the adversary can monitor the measurements of more sensors when the radius of the wormhole is larger, but in that case, it can also be detected more accurately by the NNT algorithm.

4.2 Results of the All Distances Test (ADT)

The results of the ADT algorithm are shown on Figures 7 and 8. Figure 7(a) shows the accuracy of the detection as a function of the sensors' radio range when the radius of the wormhole is 50 m, whereas Figure 7(b) shows the same

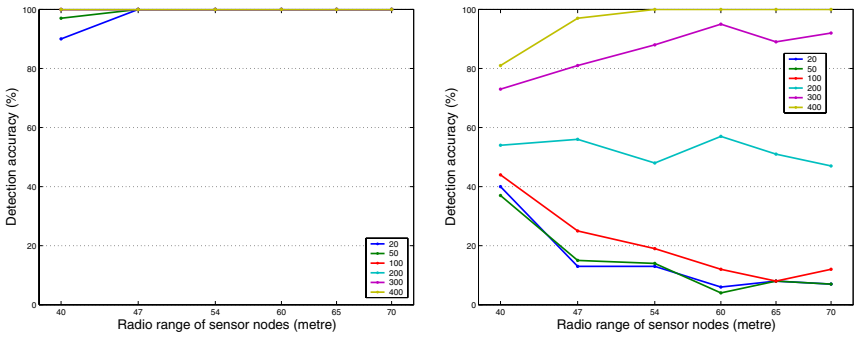

Fig. 7. Detection accuracy plotted against the radio range of the sensor nodes. The different curves belong to different distances between the areas affected by the wormhole with a radius of 50 m (a) and 16 m (b).

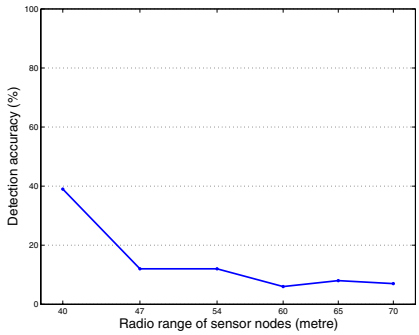

Fig. 8. Percentage of false positive wormhole detections plotted against the radio range of the sensor nodes

when the radius of the wormhole is 16 m. Similar to the NNT algorithm, the ADT algorithm performs better when the radius of the wormhole is larger. However, unlike the NNT algorithm, the ADT algorithm is not completely unusable in the case when the radius of the wormhole is 16 m. Rather, its performance depends on the distance between the areas affected by the wormhole: the higher this distance is, the more accurate the detection is. Moreover, when the distance between the affected areas is 400 m, the accuracy is close to 100% . The explanation for this is quite obvious: a longer wormhole reduces the length of the shortest paths between more distant nodes, and thus overall, it represents a larger decrease in the average length of the shortest paths between all pairs of nodes.

Regarding the percentage of the false positive alarms (Figure 8), the ADT algorithm performs quite well except for small radio ranges.

One may have expected that the detection accuracy of the ADT algorithm is independent of the radius of the wormhole. The rationale would be that no

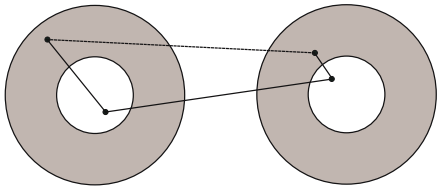

Fig. 9. Shortest paths are longer when the radius of the wormhole is smaller

matter how many new links are created by the wormhole, the important thing is that it creates shortcuts in the graph which reduce the lengths of the shortest paths between the sensors. However, this intuition is fallacious: shortest paths are indeed longer if the radius of the wormhole is smaller. As an illustration, let us consider Figure 9. The upper two nodes are directly connected if the radius is larger, whereas they are three hops away if the radius is small. This difference may seem to be small, but note that many shortest paths may use the wormhole and this two hop difference appears in each of them.

5 Conclusion and Future Work

In this paper, we have studied the problem of wormhole detection in wireless sensor networks. We proposed two mechanisms for wormhole detection that are based on hypothesis testing, and that provide probabilistic results. The first mechanism, called the Neighbor Number Test (NNT), detects the increase in the number of the neighbors of the sensors, which is due to the new links created by the wormhole in the network. The second mechanism, called the All Distances Test (ADT), detects the decrease of the lengths of the shortest paths between all pairs of sensors, which is due to the shortcut links created by the wormhole in the network. Both mechanisms assume that the sensors send their neighbor list to the base station, and it is the base station that runs the algorithms on the network graph constructed from the received neighborhood information.

We investigated the effectiveness of the two proposed mechanisms by means of simulation. Our results show that both mechanisms can detect the wormhole with high accuracy when the radius of the wormhole is comparable to the radio range of the sensors. In addition, the ADT algorithm performs better than the NNT algorithm when the radius of the wormhole is small (compared to the radio range of the sensors). In terms of false alarms, both algorithms perform reasonably well.

One disadvantage of the mechanisms that we proposed in this paper is that they detect only the presence of a wormhole, but they do not pinpoint its location. While detection is certainly the first thing that one needs to do, localization of the wormhole afterwards is also necessary for a successful defense. In the future, we intend to study if the statistical approach proposed in this paper can be extended to provide also wormhole localization services.

In this paper, we addressed the problem of wormhole detection in a static setting. In the future, we intend to extend our results to the dynamic case,

when the wormhole is not present in the system from the beginning, but it is established by the adversary during the operation of the network. To some extent, detecting a dynamic wormhole is easier than detecting a static one: if the base station detects that two sensors that previously were many hops away from each other become neighbors, then it is reasonable to assume that a wormhole has just been established between them. On the other hand, such a detection scheme would require the sensors to provide neighborhood information to the base station continuously; a prohibitive price in sensor networks. Therefore, we are interested in the trade-offs between the overhead, the speed, and the accuracy of the detection.

Acknowledgement

The work presented in this paper has partially been supported by the Hungarian Scientific Research Fund (T046664). The first author has been further supported by IKMA and by the Hungarian Ministry of Education (BÖ2003/70).

References

1. I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, E. Cayirci. Wireless sensor networks: a survey. *Computer Networks* 38:393-422, 2002.
2. I.N. Bronstein, K.A. Semendjajew, G. Musiol, and H. Muehlig. *Handbook of Mathematics*, Springer, 2004.
3. S. Čapkun, L. Buttyán, and J.-P. Hubaux. SECTOR: secure tracking of node encounters in multi-hop wireless networks. In *Proceedings of the ACM Workshop on Security in Ad Hoc and Sensor Networks (SASN)*, 2003.
4. L. Hu and D. Evans. Using directional antennas to prevent wormhole attacks. In *Proceedings of the IEEE Symposium on Network and Distributed System Security (NDSS)*, 2004.
5. Y. Hu, A. Perrig, and D. Johnson. Packet leashes: a defense against wormhole attacks in wireless ad hoc networks. In *Proceedings of the IEEE Conference on Computer Communications (Infocom)*, 2003.
6. C. Karlof and D. Wagner. Secure routing in sensor networks: attacks and countermeasures. *Ad Hoc Networks* 1:293–315, 2003.
7. R. Poovendran and L. Lazos. A graph theoretic framework for preventing the wormhole attack in wireless ad hoc networks, to appear in *ACM Wireless Networks*.
8. W. Wang and B. Bhargava. Visualization of wormholes in sensor networks. In *Proceedings of the ACM Workshop on Wireless Security (WiSe)*, 2004.