Evaluating Sinkhole Defense Techniques in RPL Networks

Kevin Weekly and Kristofer Pister
Department of Electrical Engineering and Computer Sciences
University of California Berkeley
Berkeley, California, USA
Email: {kweekly, pister}@eecs.berkeley.edu

***Abstract*—In this work, we present the results of a study on the detrimental effects of sinkhole attacks on Wireless Sensor Networks (WSNs) which employ the Routing Protocol for LLNs (Low-power and Lossy Networks). A sinkhole is a compromised node which attempts to capture traffic with the intent to drop messages, thus degrading the end-to-end delivery performance, that is, reducing the number of messages successfully delivered to their destination. The mechanism by which the sinkhole captures traffic is by advertising an attractive route to its neighbors. We evaluate two countermeasures addressing the sinkhole problem: a parent fail-over and a rank authentication technique. We show via simulation that while each technique, applied alone, does not work all that well, the combination of the two techniques significantly improves the performance of a network under attack. We also demonstrate that, with the defenses described, increasing the density of the network can combat a penetration of sinkholes nodes, without needing to identify the sinkholes.**

***Keywords*-Communication system security, Ad hoc networks, Routing protocols**

I. Introduction

The development and deployment of low-power *wireless sensor networks (WSNs)* is a growing industry fueled by the potential to replace wired infrastructure and the costs associated with installing and maintaining cabling. In achieving upwards of 99.99% reliability in their networks, commercialized instantiations are able to provide an economic and reliable WSN strategy for their clients [1].

An aspect emerging from work on WSNs is the push to standardize WSN protocols. Inspired by the ability of Internet devices to inter-operate, this push proposes that devices in WSNs should as well. There are many efforts towards this end, such as the commercial ZigBee Alliance [2] and more experimental networking stacks such as TinyOS [3], Contiki [4] and OpenWSN [5]. The latter two additionally provide open-source implementations, lowering the barrier to entry. Since open protocols ensure that any adversary is privy to the internal mechanisms of the network, robust security assurances are necessary.

For the OpenWSN project, the Internet Engineering Task Force (IETF) *Routing Protocol for Low-power and Lossy Networks (RPL)* [6] [7] is the routing protocol used and what we evaluate in this work. Specifically, we aim to investigate *sinkhole attacks* on a WSN routing layer. A sinkhole attack is a *Denial of Service (DoS)* technique employed by an internal attacker to disrupt the operation of a WSN, for instance, to sabotage an industrial process or force the owner of the network to investigate the problem. A sinkhole occurs when a compromised node performs two malicious acts: First, it attracts legitimate traffic by advertising a favorable route, e.g. through manipulation of the rank field in a *Destination Information Object (DIO)* message. Second, the sinkhole drops any legitimate data traffic routing through it, degrading the performance of the network. In contrast to previous work focusing on the prevention of attacks [8] [9], our goal is to investigate the degradation of end-to-end performance when a WSN is under attack by sinkholes.

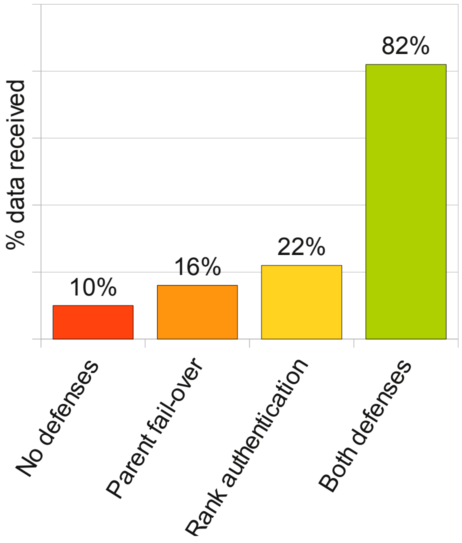

Fig. 1: End-to-end delivery ratio of simulated RPL network with 20% sinkhole penetration and 5dBm of model uncertainty. (See Section VI).

A further goal is to evaluate how two defense techniques, parent fail-over and rank authentication, allow more messages to reach their destination. The conclusive result, shown in Figure 1, demonstrates that the end-to-end message delivery ratio in an attacked WSN improves slightly with the use of two security techniques, and especially improves when both are used in tandem. Shown is a 10% delivery ratio without either defense which improves to 82% when both defenses are applied.

The paper is organized as follows: In Section II we will discuss some of the security literature related to our problem. Section III introduces the attacker model for our simulation work and Section IV describes the sinkhole defenses proposed. Section V provides important details of the simulated WSN and Section VI presents the results obtained from simulations. Section VII describes improvements to increase the efficacy of the methods and we conclude the paper in Section VIII.

II. Related Work

By design, the rank field of a DIO message must be mutable by nodes as it passes through the network. The rank field specifies how favorable a node appears to its neighbors, thus, falsifying this field is a straightforward mechanism by which a node executes a sinkhole attack. The use of one-way hash chains by the VeRA [10] and Ariadne [11] architectures offer a response to this vulnerability. Employing one of these methods ensures that any node in the network cannot falsify their topological position in the network by greater than one. Although attractive for its security benefits, a drawback of using this method is that the performance of the network degrades in a noisy environment since certain rank metrics cannot be used.

One solution proposed to defeat routing attacks more generally is the use of packet-leashes [12]. This technique uses physical properties such as node spacing and radio time of flight to verify that a multi-hop message arrives when expected. This requires an accurate propagation model and/or tight time synchronization among nodes. A visualization-based approach [13] is also proposed in which one can identify wormholes by recognizing suspicious topologies of the resulting network. This also relies on an accurate propagation model, which is difficult to construct for indoor environments.

Other solutions propose to identify compromised nodes through special nodes which eavesdrop on routing messages to verify plausibility [14]. Alternatively, all nodes have a voting mechanism to identify malicious behavior [15]. There is also a centralized approach which looks at the routing information of nodes in the network [16] to identify sinkholes. These works typically rate efficacy by false-positive and false-negative measurements, whereas the end-to-end message delivery ratio is our performance metric.

III. Threat Model

In our simulated WSN, we make the following security assumptions:

- We assume loose time synchronization of the network by the Media Access Control (MAC) layer, allowing the use of nonces for replay protection as in [17] and [18].
- We assume that all fields of DIO messages are signed by the root except for the rank and sender fields, which must be modified by the algorithm, e.g. [19].
- We assume that each non-root node is given a symmetric key shared only with the root, allowing confidentiality, integrity, and authenticity of transport layer payloads between the root and router nodes, e.g. Advanced Encryption Standard (AES) encryption [20].

The assumptions are reasonable given the technology currently available. AES encryption is already present in IEEE802.15.4 networks, as well as loose time synchronization. Asymmetric encryption is less common, but is emerging in commercial networks where security is critical. However, these assumptions alone will not prevent sinkholes from disrupting the network's operation.

The attacker is the controller of E compromised nodes $\tilde{N}_{1\ldots E}$. A compromised node has the ability to perform the following malicious behaviors:

- $\tilde{N}_i$ may drop data messages intended to be routed through it. For our simulations, $\tilde{N}_i$ always drops data packets, though future work will look into *partial sinkholes* which only selectively drop data.
- $\tilde{N}_i$ can manipulate the rank field of outgoing DIO messages, lowering it to attract more traffic. The combination of this behavior with dropping data makes the node a sinkhole. If rank verification is used, then $\tilde{N}_i$ cannot lower its true rank by more than a cost of 1.

Consider the effect of a sinkhole exhibiting these behaviors on RPL. If a low rank is broadcast by such a node, then any neighbor of the sinkhole will likely choose the sinkhole as a parent. In aggregate, this causes much of the WSN's traffic to be routed through sinkholes, severely reducing the number of packets successfully routed to their destination, the root. The problem persists even when the routing tree is reconstructed.

IV. Techniques Employed

The following two techniques have been simulated to evaluate their effect on the end-to-end delivery performance of the network:

A. Rank Verification

Rank verification ensures that compromised nodes can only lower their rank by 1. A caveat is that all edges in the graph must be weighted equally, precluding the use of finer-grained metrics. We constructed a reduced implementation of [10], which uses hash-chains to verify topological rank. Our work assumes the actual cryptographic assurances of the technique are not threatened, therefore, we do not actually perform the operations in order to speed computation.

The technique relies on a one-way hash function $h(\cdot)$, such as SHA1, and a hash chain defined by $x_{n+1} = h(x_n)$. The root begins by picking a random number x_0 and computing its hash $x_1 = h(x_0)$. The DIO rank field is augmented, or possibly replaced, with this hash when broadcast. As each legitimate node forwards the DIO message, per the normal routing tree construction, it replaces the hash value, x_i, in the message with $x_{i+1} = h(x_i)$. Sinkhole nodes transmit the hash values they received, x_i, without hashing them. Note that this is the most that a sinkhole can do to appear close to the root, since, due to the one-way nature of the hash function, it cannot determine x_{i-1}.

When choosing a parent, each node stores the hash value that the parent sent, x_p. After the routing tree has converged, then for any node, N_i,

$$\begin{aligned} p &= \hat{Rank}(N_i) \\ \hat{Rank}(N_i) &= Rank(N_i) - E_{\text{path}(i)} \end{aligned}$$

where $\hat{Rank}(N_i)$ is the perceived rank of N_i, $E_{\text{path}(i)}$ is the number of compromised nodes in the path from N_i to the root, and p is the number of hash operations performed on the original x_0 to give x_p.

After some reasonable time has passed to ensure that the routing tree has converged, the root will use a secure broadcast to distribute x_0 to all nodes in the network. Using x_0, a node can compute the required value for x_p through $\hat{Rank}(N_i)$ successive applications of $h(\cdot)$ on x_0. The expected value for x_p must match the value given by the respective parent, otherwise the child node can assume that its parent is falsifying its rank.

B. Parent fail-over

Additionally, we devised an end-to-end acknowledgment scheme, or the parent fail-over technique. This scheme adds an *unheard nodes set (UNS)* field to a DIO message when the root node first transmits it. The UNS field is signed by the root to prevent modification in transit.

The UNS is populated by node identifiers of nodes whose paths may be compromised by a sinkhole. In our simulation, non-root nodes are expected to transmit sensor data every 10 seconds. If less than 30% of these messages are received by the root, the root adds that node to the UNS. In practice, this threshold should be chosen carefully. If the threshold is too high, then natural packet loss will be interpreted as indicating a sinkhole. However, a low threshold enables a partial sinkhole to evade detection by randomly allowing messages to pass with some probability above the threshold.

A node, upon receiving a DIO with itself included the UNS, will add its parent to a local *blacklist*, effectively barring it from ever being chosen as a parent again. In practice, the blacklist should not be absolute, otherwise natural packet loss may cause a node to add all of its neighbors to the blacklist. Despite this risk, our simulations show that even in the absence of sinkholes, having a permanent blacklist results in more messages being delivered than having no blacklist at all. We attribute this to nodes blacklisting parents which are connected via low-quality links.

The parent fail-over technique has the weakness that a Sybil [21] attack may cheaply multiply the number of compromised nodes simply by having a single node pretend to be several nodes. This attack threatens several networking layers, and has some novel solutions proposed [22] [23]. We do not offer our own solution in this report, but recognize the Sybil attack can, in effect, neutralize the parent fail-over method by providing an infinite supply of candidate parents for a neighbor to choose from.

V. Procedure

A. Simulator Design

Our simulated WSN consists of 100 nodes placed randomly from a uniform distribution in a 2-dimensional square area (1km^2) and exhibits the following behaviors:

- Every 120 seconds, node 0, the root, sends a DIO message to start routing tree construction, incrementing the version each time.
- Every 10 seconds, all non-root, uncompromised nodes generate and send a data message.
- All nodes follow the semantics RPL provides, except for sinkhole nodes, which attempt rank attacks and block data packets routed through them.
- Sinkholes are chosen one-by-one from the set of uncompromised nodes so the set of sinkholes are spatially clustered.

To model packet delivery across a radio link, we use the Friis Transmission Equation [24]. Assuming we are using a 2.45GHz center frequency and commonly-available 5.6dBi antennas, the power at the receiver is given by:

$$\begin{aligned} P_r &= P_t + G_t + G_r + 20\log_{10}\left(\frac{\lambda}{4\pi R}\right) \\ &= P_t + 5.6 + 5.6 + 20\log_{10}\left(\frac{0.122\text{m}}{4\pi R}\right) \\ &= P_t - 29.06 - 20\log(R) \end{aligned}$$

where R is the distance in meters between the sender and receiver, and P_t is the transmission power (chosen to be 0dBm).

To model the uncertainty of an indoor environment we add a slow and fast-moving uncertainty to the model:

$$\begin{aligned} P_r &= P_t - 29.06 - 20\log(R) - n_{\text{slow}} + n_{\text{fast}} \\ n_{\text{slow}} &\sim \mathcal{U}(0, 40) \\ n_{\text{fast}} &\sim \mathcal{U}\left(-\frac{\text{NOISE\_BOUND}}{2}, \frac{\text{NOISE\_BOUND}}{2}\right) \end{aligned}$$

where n_{slow} is sampled at the beginning of the simulation and is known by the receiver, and n_{fast} is sampled on every message transmission and is unknown to any node.

This is a reasonable model for indoor links with high multipath effects. This transmission model boils down to a simple rule used in our simulation: a single-hop radio message is successfully transmitted if and only if

$$\begin{aligned} P_r &> -89\text{dBm} \\ -29.06 - 20\log(R) - n_{\text{slow}} + n_{\text{fast}} &> -89\text{dBm} \end{aligned}$$

where -89dBm is a reasonable *receive sensitivity* provided by current hardware.

It is important to note that a positive sample for n_{fast} could in fact be harmful for a network, as it could allow a DIO message to pass through an otherwise poor quality link. If the receiver then chooses the transmitter as its parent, subsequent data transmissions are unlikely to succeed.

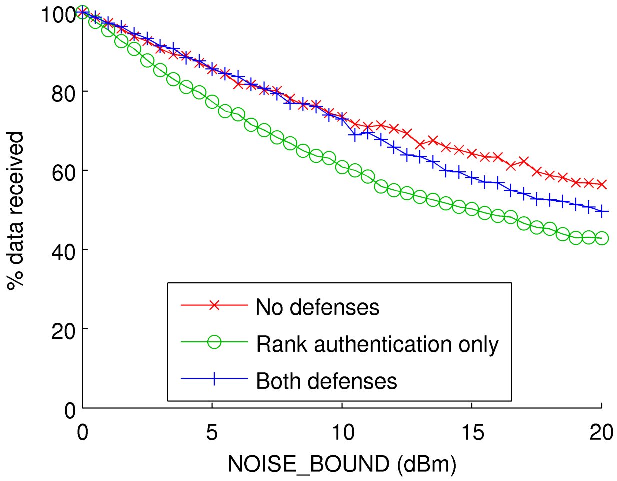

Fig. 2: Noise sensitivity analysis on an uncompromised WSN using three security configurations.

Our experiments used a custom-built discrete time RPL simulator in C++. Statistics are collected after running the simulation for 1 simulated hour to avoid measuring transient effects of network startup, and allowing the parent fail-over technique to settle. We also built tools aiding us in running our simulations: First, we constructed a visualization tool in JAVA to read trace files from the simulator, display nodes and messages in-flight, and show the routing tree structure as it changes over time. Second, we wrote a script in Python to automate testing by writing parameters to a header file, compiling the C++ code, and running it in a new process.

VI. Results

In this section, we present the results of experiments run during this study. The primary metric used is end-to-end delivery ratio, calculated by:

$$(\text{messages received at root}) \div (\text{messages sent})$$

Although simulated messages are generated by compromised nodes, they are *not* counted in the statistics.

Figure 2 shows the results of an experiment conducted to determine if there is a negative impact of the two security techniques on an uncompromised network. Noticing the detrimental effect of rank verification, we hypothesize that it is due to the requirement of equal edge weights in the connectivity graph. In the presence of uncertainty, ranking potential parents by link quality is a beneficial technique. Fortunately, as the results show, the parent fail-over technique naturally introduces a feedback mechanism blacklisting parents with low link quality.

For the next experiment, we ensured that our results were not biased by the choice of NOISE_BOUND. Intuition tells us that the parent fail-over technique could be especially sensitive to dropped packets, due to the permanent nature of the parent blacklist. In Figure 3, we show the results of increasing NOISE_BOUND on a WSN employing both defenses. Observing that the results in Figure 2 show a similar shaped curve, we hypothesize that the defenses do not multiply the detrimental effect of noise. Moreover, because the “80 sinkholes” curve is relatively flat, we observe that the detrimental effect of a large number of sinkholes in the network overshadows the effect of uncertainty. This is also an observation found in Figure 5.

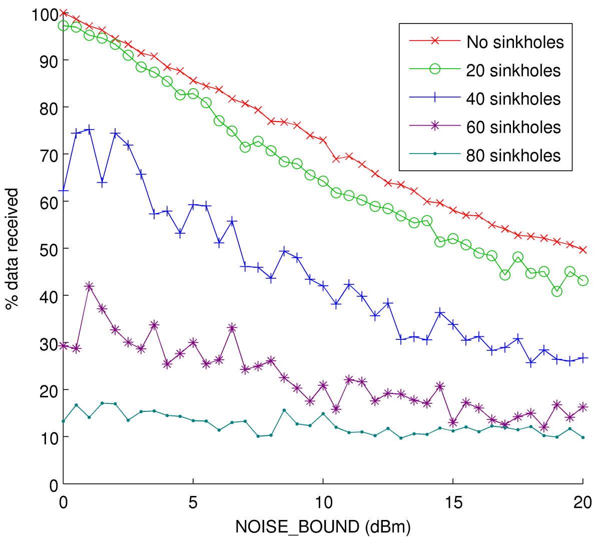

Fig. 3: Noise sensitivity analysis on a compromised WSN using both parent fail-over and rank authentication defenses.

Figure 4 shows the detrimental effect of sinkholes on a network as the number of them grows. Note that the maximum number of sinkholes tested is 98, since there must be at least one legitimate node and one root in the network for the calculation of end-to-end delivery ratio. The control experiment (i.e. “no defenses”) is shown to have the least successful transmissions, followed by the parent fail-over and rank authentication methods applied alone. We demonstrate

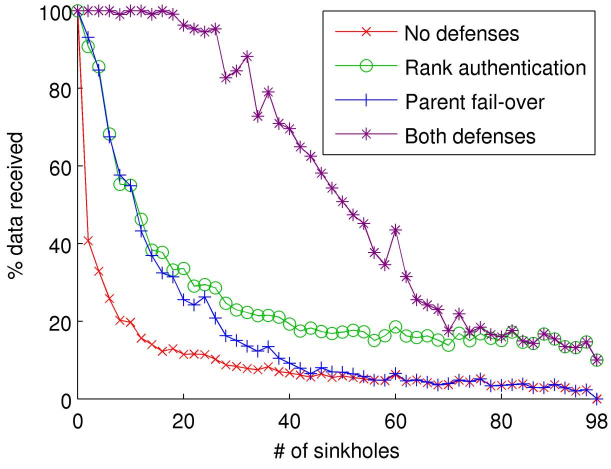

Fig. 4: Analysis of attacker penetration on successful data transmission using four different security strategies. NOISE_BOUND=0.

that the highest delivery ratio occurs when rank authentication is combined with the parent fail-over method. Our hypothesis is that the two techniques are synergistic in the following way: Using only parent fail-over, each router will have a large number of neighboring sinkholes advertising a rank of 0. Thus, these victim routers may never choose a legitimate node as a parent. Rank authentication addresses this by preventing neighboring sinkholes from advertising a rank of 0, however there is no feedback if a node chooses a sinkhole as a parent. Using both techniques concurrently, each technique mitigates the others' weakness. It is also interesting to observe that, at 98 sinkholes, there is still some degree of success for the single remaining legitimate node to find the root as its parent.

In the next experiment, we add a 5dBm NOISE_BOUND to model packet loss of a realistic deployment and show the results in Figure 5. The amount of uncertainty was chosen to be significant, but not enough to cripple the network and eclipse the effect of the sinkholes. We can see that the shape of the results look similar to Figure 4, albiet with overall worse performance. This indicates our deductions from the previous case remain valid. An interesting discovery is that, for a low number of sinkholes, the loss seems to be driven by the uncertainty, but the uncertainty starts to become irrelevant when the number of sinkholes is over 50.

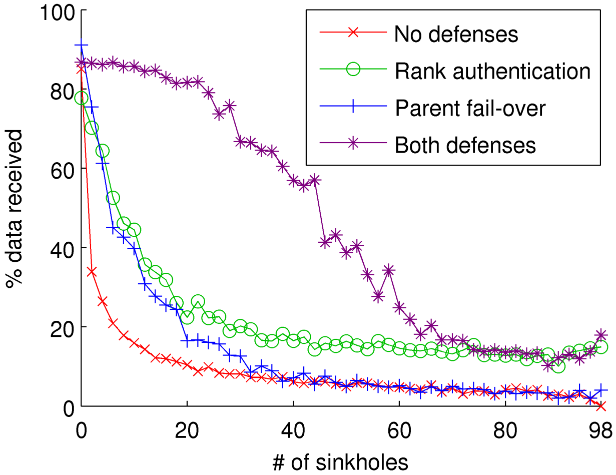

Fig. 5: Analysis of attacker penetration effect on successful data collection in the presence of uncertainty (NOISE_BOUND = 5dBm).

For the final experiment of this paper we studied how a WSN employing the described defensive techniques responds to the average topological height of the network. To incite a larger average number of hops a data message must travel to reach the root, the spatial size of the network was increased. We show the results in Figure 6, plotting both the end-to-end success rate and the average number of hops for a successful transmission. In the "no sinkholes" case, we see that the *breaking point* of an uncompromised network is around 3km × 3km (7 hops), after which too many nodes are outside the reliable transmission range of the network. As we increase the number of sinkholes, the breaking point moves lower, such that for the "80 sinkholes" case, we see significant performance degradation start at 100m × 100m, or 1.1% of the area of the uncompromised network. This result shows that, in a WSN, using the two security techniques described, one could counter the effects of sinkholes by adding more router nodes, thus increasing the spatial density of the network. For low numbers of sinkholes, this may be an economical option

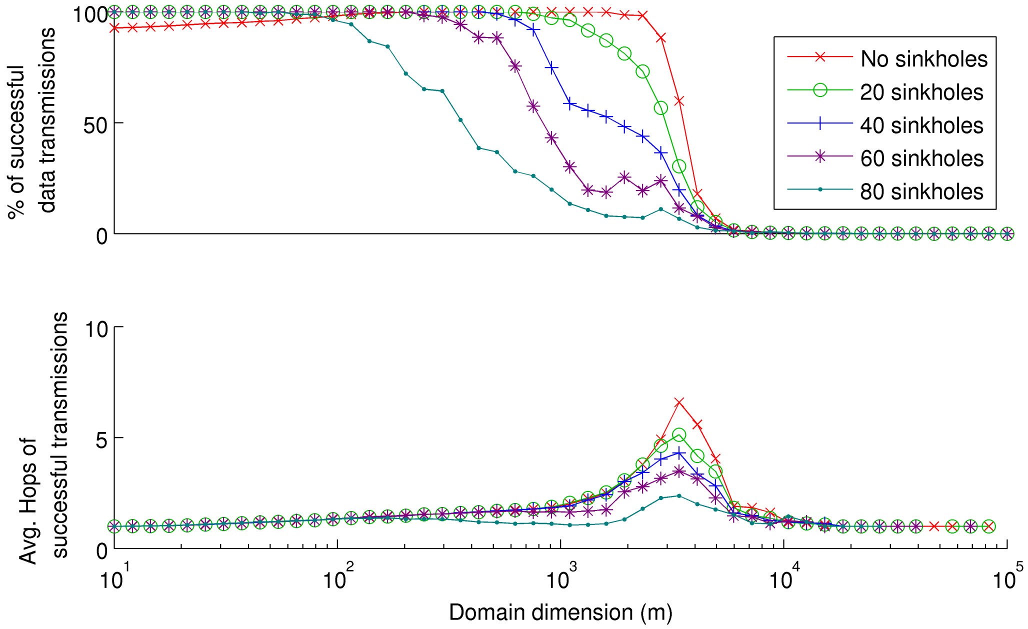

Fig. 6: Analysis of domain size on routing performance and number of hops in network. NOISE_BOUND=0.

to restore network operation. It is an intuitive result, but it is helpful to verify the effect.

VII. Extensions and Future Work

We envision the following extensions that could increase the efficacy of the techniques:

- First, the parent fail-over technique creates an inefficiency which may cause many legitimate nodes to be blacklisted. Consider a single sinkhole blocking messages from its entire subtree. On the next DIO update, the UNS will contain all of the nodes in the attacker's subtree. Each will consequently place their parents into their blacklists. We propose that a node, N_i, should instead switch only after $Rank(N_i)$ DIO messages have included N_i in the UNS, as this would give the node's ancestors the chance to fail-over first.
- Second, a realistic implementation of RPL includes multiple routing trees operating concurrently. We could exploit the redundancy to better identify sinkholes. For example, if two paths are dropping messages and share a single router along them, that router is likely to be a sinkhole. A statistical approach would be appropriate to integrate these types of observations.
- Finally, the logging information that nodes and roots collect could assist in locating sinkhole nodes. The challenge is that compromised nodes are likely to falsify the logs to blame their peers for malicious behavior.

Looking forward, we believe this work opens several avenues of research into mitigating the effects of attacks on WSNs. We are interested in implementing other strategies such as parent diversity, mesh networking, and two-hop verification (i.e. a child monitors its parent's transmissions), investigating their benefits as well. There is the problem of increasing performance in a lossy environment which we believe shares common solutions with the sinkhole problem. Additionally, we would like to record metrics as communication bandwidth and energy overhead in the simulation. Finally, measuring a physical implementation on OpenWSN hardware will validate these simulated results.

VIII. Conclusion

Reliability in an adversarial environment is a developing topic important for the acceptance of WSNs as a replacement for traditionally wired networks. In this paper, we presented the results of applying two security techniques on a simulated WSN under attack by sinkholes. We demonstrated that the two techniques, parent fail-over and rank authentication, mitigate the detrimental effects of sinkholes on an RPL network. Furthermore, the combination of the two techniques is significantly more effective than either technique alone. Finally, we showed that a network under attack by sinkholes and employing these defenses can regain high end-to-end performance by increasing the density of routers in the network.

Acknowledgment

Special thanks to Vern Paxon for his guidance in shaping the project.

References

[1] L. Doherty and D. Teasdale, "Towards 100% reliability in wireless monitoring networks," in *Proceedings of the 3rd ACM International Workshop on Performance Evaluation of Wireless Ad Hoc, Sensor and Ubiquitous Networks*. ACM, 2006, pp. 132–135.

[2] (2012, May) Zigbee alliance. [Online]. Available: http://www.zigbee.org/

[3] (2012, May) Tinyos home page. [Online]. Available: http://www.tinyos.net/

[4] A. Dunkels *et al.*, "Contiki-a lightweight and flexible operating system for tiny networked sensors," in *29th Annual IEEE International Conference on Local Computer Networks*. IEEE, 2004, pp. 455–462.

[5] T. Watteyne *et al.*, "OpenWSN: A standards-based low-power wireless development environment," *Transactions on Emerging Telecommunications Technologies*, in-press.

[6] T. Winter *et al.*, "RPL: IPv6 routing protocol for low power and lossy networks," Mar. 2011. [Online]. Available: http://tools.ietf.org/html/draft-ietf-roll-rpl-19

[7] J. Vasseur *et al.*, "RPL: The IP routing protocol designed for low power and lossy networks," *Internet Protocol for Smart Objects Alliance*, 2011. [Online]. Available: http://www.cs.berkeley.edu/~jwhui/6lowpan/IPSO-WP-7.pdf

[8] T. Tsao *et al.*, "A security framework for routing over low power and lossy networks," Mar. 2010. [Online]. Available: http://tools.ietf.org/html/draft-tsao-roll-security-framework-02

[9] C. Karlof and D. Wagner, "Secure routing in wireless sensor networks: attacks and countermeasures," in *Proceedings of the 1st IEEE International Workshop on Sensor Network Protocols and Applications*, may 2003, pp. 113 – 127.

[10] A. Dvir *et al.*, "VeRA - version number and rank authentication in RPL," in *IEEE 8th International Conference on Mobile Adhoc and Sensor Systems*, oct. 2011, pp. 709 –714.

[11] Y.-C. Hu *et al.*, "Ariadne: A secure on-demand routing protocol for ad hoc networks," in *Proceedings of the 8th ACM International Conference on Mobile Computing and Networking*, Sep. 2002.

[12] Y. Hu *et al.*, "Wormhole attacks in wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 2, pp. 370–380, 2006.

[13] W. Wang and B. Bhargava, "Visualization of wormholes in sensor networks," in *Proceedings of the 3rd ACM Workshop on Wireless Security*. ACM, 2004, pp. 51–60.

[14] A. Le *et al.*, "Specification-based IDS for securing RPL from topology attacks," in *Wireless Days (WD), 2011 IFIP*, oct. 2011, pp. 1 –3.

[15] I. Krontiris *et al.*, "Intrusion detection of sinkhole attacks in wireless sensor networks," in *Proceedings of the 3rd International Conference on Algorithmic Aspects of Wireless Sensor Networks*. Springer-Verlag, 2007, pp. 150–161.

[16] E. Ngai *et al.*, "On the intruder detection for sinkhole attack in wireless sensor networks," in *IEEE International Conference on Communications*, vol. 8. IEEE, 2006, pp. 3383–3389.

[17] T. Watteyne *et al.*, "Reliability through frequency diversity: why channel hopping makes sense," in *Proceedings of the 6th ACM symposium on Performance evaluation of wireless ad hoc, sensor, and ubiquitous networks*. ACM, 2009, pp. 116–123.

[18] K. Pister and L. Doherty, "Tsmp: Time synchronized mesh protocol," *IASTED Distributed Sensor Networks*, pp. 391–398, 2008.

[19] M. Luk *et al.*, "Seven cardinal properties of sensor network broadcast authentication," in *Proceedings of the 4th ACM Workshop on Security of Ad Hoc and Sensor Networks*, ser. SASN '06. New York, NY, USA: ACM, 2006, pp. 147–156. [Online]. Available: http://doi.acm.org/10.1145/1180345.1180364

[20] N. Sastry and D. Wagner, "Security considerations for ieee 802.15.4 networks," in *Proceedings of the 3rd ACM Workshop on Wireless Security*. ACM, 2004, pp. 32–42.

[21] J. Douceur, "The sybil attack," *Peer-to-Peer Systems*, pp. 251–260, 2002.

[22] B. Levine *et al.*, "A survey of solutions to the sybil attack," *University of Massachusetts Amherst, Amherst, MA*, 2006.

[23] J. Newsome *et al.*, "The sybil attack in sensor networks: analysis & defenses," in *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks*. ACM, 2004, pp. 259–268.

[24] H. Friis, "A note on a simple transmission formula," *proc. IRE*, vol. 34, no. 5, pp. 254–256, 1946.