

SIGF: A Family of Configurable, Secure Routing Protocols for Wireless Sensor Networks

Anthony D. Wood, Lei Fang,
John A. Stankovic
Department of Computer Science
University of Virginia

{wood, leifang, stankovic}@cs.virginia.edu

Tian He
Dept. of Computer Science and Engineering
University of Minnesota
tianhe@cs.umn.edu

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols—Routing Protocols

General Terms: Security, Design, Performance

Keywords: Wireless Sensor Network, Secure Routing

ABSTRACT

As sensor networks are deployed in adversarial environments and used for critical applications such as battlefield surveillance and medical monitoring, security weaknesses become a big concern. The severe resource constraints of WSNs give rise to the need for *resource bound security* solutions.

In this paper we present SIGF (Secure Implicit Geographic Forwarding), a configurable secure routing protocol family for wireless sensor networks that provides “good enough” security and high performance. By avoiding or limiting shared state, the protocols prevent many common attacks against routing, and contain others to the local neighborhood.

SIGF makes explicit the tradeoff between security provided and state which must be stored and maintained. It comprises three protocols, each forming a basis for the next: SIGF-0 keeps no state, but provides probabilistic defenses; SIGF-1 uses local history and reputation to protect against certain attacks; and SIGF-2 uses neighborhood-shared state to provide stronger security guarantees.

Our performance evaluation shows that SIGF achieves high packet delivery ratios with low overhead and end-to-end delay. We evaluate the security of SIGF protocols under various security attacks and show that it effectively contains the damage from compromised nodes and defends against black hole, selective forwarding, Sybil, and some denial of service attacks.

1. INTRODUCTION

Security is critical for many wireless sensor network applications such as battlefield surveillance, medical monitoring,

and emergency response. However, many security mechanisms developed for the Internet or ad-hoc networks cannot be applied directly to wireless sensor networks (WSNs) due to their limited resources in computation, memory, communication bandwidth, and energy.

The severe resource constraints of WSNs give rise to the need for *resource bound security* solutions. There are at least two interesting aspects of this concept. First, individual security mechanisms must be efficient in memory, computation, energy and bandwidth. For example, certain cryptographic schemes are inappropriate because ciphertext message expansion results in costly memory, bandwidth and energy use. Second, the resource consumption of all security mechanisms installed together at a node must not exceed the amount of resources allocated for security and they cannot degrade performance to an unacceptable level during normal operation nor when an attack is underway.

It is not possible in today’s state of the art to include strong security mechanisms for *each* of the services at a node such as medium access control, routing, localization, time synchronization, power management, sensing, and group management. Consequently, even if a secure (to a wide variety of attacks) routing protocol is implemented, it may suffer from low efficiency and would not protect against attacks on the other services.

Our approach for resource bound security is to have minimal *active* security protection. This results in very high performance and minimal resource consumption when no attacks are underway. Then upon detecting an attack or if the system designers expect increased threats, the appropriate security mechanism is activated. The result is not 100% security protection—but *good enough* security, activated at the right time. This general approach makes it possible to have high performance and to react to current security attacks, and is even more evolvable to new attacks than approaches that fix a set of solutions into a node.

In this paper we present Secure Implicit Geographic Forwarding (SIGF), a family of configurable secure routing protocols that follow this general approach. For a complete WSN solution similar families of protocols would be required for each of the other services.

SIGF is based on IGF [1], a nondeterministic Network/MAC hybrid routing protocol that is completely stateless. This allows it to handle network dynamics effortlessly, and intrinsically limits the effects of a compromised node to a local area. There are no routing tables to corrupt, since forwarding decisions are made as late as possible—when a packet is ready to transmit over the air. Nevertheless, it

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SASN’06, October 30, 2006, Alexandria, Virginia, USA.
Copyright 2006 ACM 1-59593-554-1/06/0010 ...\$5.00.

is susceptible in the local neighborhood to a simple CTS rushing attack [9].

SIGF comprises three protocols which extend IGF and populate the gap between pure statelessness and traditional shared-state security. SIGF-0 keeps no state, but uses non-determinism and candidate sampling to achieve high packet delivery ratios probabilistically. SIGF-1 keeps local state, building reputations for its neighbors to aid in next-hop selection. SIGF-2 uses state shared with neighbors to provide the strongest defense against attack, yet at the greatest cost. Each protocol encompasses the features of the previous, layering additional mechanisms to defend against more sophisticated attacks. The layered family of protocols allows a network to activate only the protections currently necessary, and to change to stronger ones only if they are warranted.

We evaluate the performance of each protocol by simulating with no attacks, and with black hole, selective forwarding, Sybil, and denial of service attacks. We show that each protocol represents a tradeoff between state and security, and that despite keeping no state, SIGF-0 performs well.

We make several contributions in this work. First, we show that even with no security countermeasures, the base protocol IGF has desirable attack containment properties, but nevertheless falls to several attacks that completely disrupt communication in a local neighborhood. Then we present the design and evaluation of SIGF, a secure routing protocol family built upon IGF. In our tested scenarios, under black hole attack from a single node the stateless SIGF-0 protocol maintains 45% packet delivery ratio (PDR), while the reputation-based SIGF-1 achieves 97% PDR. In the six-node Sybil attack, SIGF-1 still delivers 83% of packets. To the best of our knowledge, SIGF is the first configurable routing protocol for WSNs that makes explicit the tradeoffs between resources required and security provided, and enables resource bound security that is both efficient and effective.

It is possible that some WSNs require much stronger security than what our dynamic approach offers. However, no perfect solution exists—nor is likely to exist on severely resource constrained devices. Our approach, as exemplified by the SIGF routing protocols presented in this paper, can significantly improve security (not make it absolute), allow operation in the presence of attacks, and support a high performance system.

For example, consider a military surveillance, tracking, and classification system for rare-event areas. The system may operate undetected for months, with no attacks and great performance using lightweight protocols. When an adversary eventually detects the system and drops in a node to create a routing black hole, then our solution would activate a form of secure routing that can recover from this attack. In most cases the loss of some messages is not critical because of the periodic and highly redundant nature of the system.

The rest of this paper is organized as follows. In the next section we briefly review our foundational routing protocol, IGF. Then Section 3 describes system assumptions and routing attacks on IGF. Section 4 presents SIGF, our secure routing protocol family. In Section 5 we present our experiments and a detailed evaluation of the protocols under various attacks. Finally, we discuss related work in Section 6 and then conclude.

2. IMPLICIT GEOGRAPHIC FORWARDING

Our foundational routing protocol is Implicit Geographic Forwarding (IGF) [1], which is completely stateless, without dependence on knowledge of the network topology or the presence/absence of any other nodes. It makes non-deterministic routing decisions, implicitly allowing receiving nodes to determine a packet’s next-hop at transmission time. IGF couples the routing and MAC components into a single integrated Network/MAC protocol. It identifies the best forwarding candidate during MAC-layer handshaking at the instant a packet is sent.




Figure 1: Forwarding area for message sender S .

Figure 1 presents an example topology, where source node S transmits a message toward D . Routing is integrated with the RTS/CTS hand-shake of MACA/802.11 [10] DCF MAC protocols.

The communication handshake for this example topology is shown in Figure 2. It begins when the sender S ’s NAV timer is zero and it carrier senses an idle channel for DIFS time. Having verified that the channel is free, S broadcasts an Open RTS (ORTS) containing its location S and destination D .




Figure 2: IGF handshake timeline.

Neighbors are eligible to forward the message if they are within a 60° sextant centered on the direct line from the sender to the destination (the forwarding area). We call the nodes in the forwarding area *candidate nodes*. Such nodes set a CTS Response timer inversely proportional to a weighted sum of their distance from the sender, remaining energy, and perpendicular distance to a line from the sender to the destination. This favors the nodes that are more desirable for forwarding.

On the expiry of a node’s CTS response timer, it responds with a CTS packet, and the data is transferred from the Open RTS sender in a DATA message. The valid duration of the CTS timers is called the CTS-response window. Ideally, other candidate nodes can hear the CTS (by virtue of lying inside the sextant) and cancel their timers before the end of the window. Therefore, in IGF, a single node with the shortest CTS response timer responds to the ORTS.¹

¹IGF deals with network voids by shifting the forwarding

Since IGF keeps no routing state information, it provides fault tolerance and is robust under topology changes. It also eliminates expensive communication for routing and neighbor information maintenance, and the associated routing latency. Using the concept of *lazy binding*, the IGF protocol defers next-hop selection until the packet forwarding operation actually happens at a sending node. Lazy binding dramatically reduces the chance that packets are forwarded to a node that fails, moves out of range, or transits to a sleep state, and also enables the use of recently awakened or newly arriving nodes.

To avoid neighbor tables and do lazy binding, IGF must use an RTS/CTS handshake. This overhead is well paid back in performance, however, as shown in the experimental results.

Compared with other established protocols for sensor and ad-hoc networks (such as GPSR [14], DSR [11], and LAR [16]), IGF achieves up to a ten-fold increase in delivery ratio and significantly reduces both end-to-end delay and control overhead. It is therefore a good protocol to serve as a foundation for secure routing.

IGF has no routing tables, so it naturally confines the attacker's impact to the neighborhood and prevents attacks that spoof, alter, or replace routing information. This is a significant advantage over link-state and distance vector routing protocols, which must carefully manage updates and route requests to avoid contamination by attackers.

Unfortunately, a single attacker can completely disrupt routing for all of its neighbors. For example, the attacker arranges for itself to be chosen as the next-hop relay simply by sending an immediate CTS message upon receiving an ORTS. When the attacker gets the DATA, it replies with an ACK, but drops the DATA packet. The packet delivery ratio becomes zero—a simple attack, but devastating in the local neighborhood.

We designed SIGF to secure routing in the local neighborhood while preserving the performance and attack containment properties of IGF.

3. ASSUMPTIONS AND ATTACKS

Routing is an essential service for enabling communication in sensor networks, and is therefore potentially the target of many different attacks. First, we identify our assumptions about the system. We review the general classes of attacks on sensor network routing, then focus on attack mechanisms specific to our protocol in the next section.

3.1 System Assumptions

We assume that radio links are insecure, i.e., attackers may eavesdrop on radio transmissions, inject messages, and record and later replay messages. If an attacker is able to interact with the routing protocol, it can also drop messages for which it is responsible. Attackers possess hardware capabilities similar to that of legitimate nodes, and wireless transmissions use the same power levels.

Network nodes move only infrequently or slowly once deployed, and know their own locations. They may additionally know that of their neighbors (in SIGF-1 and SIGF-2). Nodes know the location of important resources, like base stations, and use it for geographic routing.

sextant to the side and retrying [1]. SIGF inherits this mechanism, which we do not discuss further in this paper.

We do not require time synchronization among nodes. For SIGF-0 and SIGF-1, no shared keys are required between nodes in the network. SIGF-2 assumes the presence of pairwise-shared keys in the neighborhood, which may be fulfilled by many different key distribution schemes in the literature [4, 2, 27]. Nodes trust their own clocks, measurements, and storage.

3.2 Routing Attacks

Karlof and Wagner [13] and others [24, 20] have systematically studied attacks on routing protocols. We summarize these attacks below, noting whether they are applicable to IGF (and therefore to SIGF). Then we discuss those attacks which are not obviously thwarted in greater detail.

1. *Routing state corruption.* By spoofing, altering, or replacing routing information, attackers are able to create routing loops, attract or redirect network traffic, increase end-to-end delay, etc. IGF keeps no information, and SIGF keeps only locally generated information.
2. *Wormholes.* In this attack, an adversary tunnels messages received in one part of the network over a low latency link and replays them elsewhere. Since IGF chooses the next-hop dynamically, a wormhole does not cause disruption when it ceases to operate.
3. *HELLO floods.* An attacker convinces nodes in the network that the attacker is a neighbor by broadcasting HELLO messages with high energy. As with the wormhole attack, dynamic routing in IGF prevents disruption by a HELLO flood.
4. *Black holes.* In a black hole attack, an adversary or compromised node lures nearly all the traffic from a particular area through itself, where the messages are dropped. We further discuss this attack below.
5. *Selective forwarding.* Attackers selectively forward packets instead of faithfully forwarding all received packets or completely dropping all packets. At one end of the spectrum, messages are rarely dropped. At the other end is a black hole attack. We group this attack with the black hole attack since its mechanism is the same and consider its impact on IGF.
6. *Sybil attack.* In the Sybil attack, a malicious node behaves as if it were a larger number of nodes by impersonating other nodes or simply by claiming false identities. We further discuss this attack below.
7. *Denial of Service.* Most attacks result in a denial of service of some sort, but this moniker is usually reserved for attacks that waste resources or disrupt service in a way that far exceeds the effort required of an attacker. Message amplification and jamming are general examples. We consider specific mechanisms for mounting this attack on IGF below.

In an insider attack, a compromised node uses any means available to legitimate nodes to disrupt the protocol or perform a specific attack listed above. All state, including keys possessed by the node, may be used by the attacker.

Since IGF keeps no routing tables, it prevents *a priori* attacks such as state corruption, wormholes, and HELLO floods. Further, the impact of all attacks is limited to the local area, since routing is fully distributed and independent from hop to hop. IGF and SIGF do not trust neighboring nodes to behave correctly, so they are resistant to attacks from outsiders and insiders alike.

The main attacks available to an adversary are to create a black hole, pose as multiple identities (Sybil attack), or disrupt the routing protocol through denial of service attacks. We describe specific mechanisms for performing these attacks on IGF in the next sections. When we describe and evaluate SIGF in Sections 4–5, we focus particularly on its resilience to these attacks.

3.2.1 Black Hole / Selective Forwarding Attack

Within the local neighborhood, the easiest way for an attacker to create a black hole is to manage to always be selected by neighbors as the next hop, whether this is proper, or not.

In the *CTS rushing attack*, an attacker exploits the cooperative nature of IGF’s next-hop selection. When an Open RTS (ORTS) message is received, neighbors set timers proportional to their desirability as forwarding candidates. The attacker disregards this mechanism and always replies immediately with a CTS, volunteering to forward the packet. Once selected as the next relay, the attacker may modify, totally drop (black hole attack) or selectively forward the DATA message. This attack is very effective against IGF, easy to perform, and requires moderate power consumption, as it is completely reactive.




Figure 3: CTS Rushing Attack by A against S.

Figure 3 shows how this attack works. When attacker A overhears an ORTS message, it sends a CTS message, whether it is in the forwarding area or not. Other nodes overhear the CTS from the attacker and abort the protocol. Unsuspecting ORTS senders in the neighborhood of the attacker always choose to send their messages into the black hole created by A.

3.2.2 Sybil Attack

In a Sybil attack, an attacker illegitimately claims to be multiple nodes by sending messages with different identities and locations. Its additional identities are virtual Sybil nodes. Without cryptographic authentication, a receiver of a message cannot determine the true identity of its originator, and does not know how many of the claimed identities are truly unique. Our foundational routing protocol IGF is vulnerable to Sybil attack because it does not maintain any neighborhood state with which to validate the identities.

Identity and Location. A Sybil node can either fabricate a new identity, or steal an identity from a legitimate node [20]. In our experiments, an attacker creates several Sybil nodes surrounding its true location and assigns each either a random or fixed location.

Communication. We assume Sybil nodes can communicate directly with legitimate nodes in the following way.

When a legitimate node sends a message to a Sybil node, the attacker overhears the message. Likewise, messages sent from Sybil nodes are actually from the attacker, but with the proper identity enclosed.




Figure 4: Node A performs a Sybil attack against S.

Communication with a Sybil node is illustrated in Figure 4. After receiving an ORTS message, the attacker sends a CTS addressed from one of the Sybil nodes. Once the Sybil node is selected as the next relay, the attacker overhears and acknowledges the DATA. It can then drop, tamper, or forward the DATA in a black hole or selective forwarding attack.

3.2.3 Denial of Service Attack

The goal of this type of attack is to deny service to the nearby nodes in a manner that is less intrusive and costly than jamming. The attacker partially executes the IGF protocol to cause nearby nodes to waste energy transmitting messages, waste time waiting on completion of the protocol, or prematurely abort the protocol. We describe two specific attacks which cause denial of service by recording and replaying legitimate messages.

In an *ORTS replay* attack, a node captures an overheard ORTS message and subsequently replays it repeatedly. Each time it is replayed, neighbors of the attacker respond with CTS messages and wait for data exchange. The wireless channel cannot be used in this local neighborhood for legitimate traffic during the CTS collection window.

In a *CTS replay* attack, the old CTS message falsely causes other eligible receivers in IGF to abort the protocol (cancel their CTS response timers). The ORTS sender selects an unsuspecting or absent node (the originator of the captured CTS) as the next hop. The sender transmits the DATA, wasting energy and channel capacity, and then must retry or drop the DATA message when no acknowledgement is forthcoming. A captured ACK could be replayed by the attacker as well, causing the sender to believe the transmission was successful.

This attack is less costly to the attacker than an ORTS replay because it is reactive: the protocol is only disrupted when a neighbor actually tries to send a message.

4. SIGF: SECURE IGF

We propose a novel secure routing protocol family, called Secure IGF (SIGF) which keeps the advantages of dynamic binding in IGF, yet provides effective defenses against the attacks discussed above. The protocols provide tradeoffs between security and state maintenance, and configurability that can be adapted at runtime.

```

1 if (include destination)
2   ORTS  $\leftarrow \langle S, S_{location}, D, D_{location}, FwdArea \rangle$ 
3 else
4   ORTS  $\leftarrow \langle S \rangle$ 
6 broadcast ORTS message
8 /* Every neighbor N receives ORTS message, and if
   in FwdArea, sets CTS response timer
   proportional to next-hop desirability, sending
   CTS =  $\langle N, N_{location} \rangle$  upon expiry. */
10  $CTS_{candidates} \leftarrow \emptyset$ 
11 while (collection window open)
12   if (CTS received AND  $N_{location} \in FwdArea$ )
13     add N to  $CTS_{candidates}$ 
15 choose  $C \in CTS_{candidates}$  for next hop
16 send DATA to C

```

Algorithm 1: SIGF-0 next-hop selection for message from current node S to ultimate destination D .

The configurability of the SIGF protocol family gives a significant advantage over other more static routing protocols. Some provide no security, while others provide strong guarantees—but at the cost of more assumptions, computation, and communication. These higher costs must be borne even when no attacks are occurring. SIGF protocols can be selected and configured for the security requirements of a particular deployment.

Network planners can select among three classes of security solutions, grouped by the amount of state they keep: no state at all (SIGF-0), locally generated state (SIGF-1), and pairwise-shared state within the neighborhood (SIGF-2). This choice is currently static, but in the future it will be dynamically adjustable.

SIGF-0 is a stateless protocol that maintains no routing information, but provides only probabilistic defenses against attack. SIGF-1 keeps limited information learned from interactions with neighbors. SIGF-2 uses keys and sequence numbers shared among neighbors to provide cryptographic guarantees in routing. Each protocol is a subset of the next. That is, SIGF-1 uses mechanisms from SIGF-0, and likewise SIGF-2 uses some from SIGF-1.

The main weakness of a last-instant dynamic binding approach, as used by IGF, is in the selection of the next-hop relay. Each of these protocols uses different means to prevent or minimize the probability of selecting an attacker as the relay, while achieving high packet delivery rates with low delay and overhead.

In the following sections we present each protocol in turn.

4.1 SIGF-0: Stateless Secure IGF

SIGF-0 is the basis of the other protocols in the Secure IGF family. Without keeping forwarding history or information about neighbors, it chooses the next-hop relay non-deterministically and dynamically. This lessens, but does not eliminate, the chance of selecting an attacker in the neighborhood.

The logic for sending a message from source S to des-

tinuation D is shown in Algorithm 1. The ORTS message (as described in Section 2) is constructed in Lines 1–4 and broadcast to the one-hop neighbors in Line 6. Neighbors of S that receive the ORTS and which are in the forwarding area start CTS response timers. Upon timer expiry, a node sends a CTS response that includes its own location. In Lines 10–13, node S collects CTS responses until the collection window closes. Then a candidate C is chosen among the responders and the DATA is relayed to node C .

The algorithm is configurable in four dimensions, each of which is described here. Each is annotated with the list of options and the line number in Algorithm 1 where it appears.


1. **Forwarding Area** $\in \{60^\circ \text{ sextant, closer, whole neighborhood}\}$ *Line 2*

In the foundational routing protocol IGF, a 60° sextant toward the destination is always used as the forwarding area. This gives some assurance that CTS responders can overhear each other and cancel their timers.

In the presence of multiple neighboring adversaries, however, this sextant may not provide enough responses from which to select. Low-density deployments allow attackers to fill the CTS candidate set to the exclusion of legitimate forwarders.

SIGF-0 allows the use of

larger forwarding areas, since for a given number of attackers this increases the probability of selecting a legitimate node. In addition to the 60° sextant, any node that is closer to the destination than the sender may respond, or all neighbors may respond. This does not increase the number of attackers included, since they already may disregard the forwarding area when attacking.



Performance is affected both by allowing messages to take longer paths, and by lengthening the collection window to accommodate greater CTS candidates. However, this is offset by the ease with which multiple attackers may capture forwarding when the narrower sextant is used. Allowing more neighbors to be considered in the forwarding area does not automatically cause worse performance when there are no attacks, since correct nodes still respond according to their desirability for forwarding (as described in Section 2).

2. **Collection Window** $\in \{\text{one responder, fixed multiple, dynamically lengthened}\}$ *Line 11*

SIGF-0 collects one or more CTS responses before choosing the next-hop relay among them. IGF closes the collection window immediately upon receiving the first CTS, but this is vulnerable to the CTS rushing attack presented earlier. The attacker disregards the correct response delay and responds first, creating a black hole in the neighborhood. Still, this option is available in SIGF-0 since it provides best performance (lowest delay and overhead) when no attacker is present.

By allowing a longer collection window, SIGF-0 collects more CTS messages before selecting a relay. The ORTS sender waits a fixed amount of time, storing

CTS responses. One is chosen according to the criteria given in the next part. A fixed-length window gives predictability and constant cost, and allows CTS response timers to be scaled *a priori* to avoid unnecessary contention during the window. A flag is included in the ORTS to prevent CTS responders from aborting the protocol when another CTS is overheard.

If not enough CTS responses are received, the window may optionally be extended dynamically. At a greater cost in delay, this allows the ORTS sender to collect enough responses to give better assurance that an attacker is not chosen.

3. **Forwarding Candidate Choice** \in {first, by priority, random, multiple} *Line 15*

Given a set of forwarding candidates collected during the window ($CTS_{candidates}$ in Algorithm 1), this parameter determines how one is chosen to be the next-hop relay. IGF always chooses the first responder, which is vulnerable to the CTS rushing attack. We allow this option since it is compatible with IGF and because it is most efficient when no attackers are present.

Selecting by priority means choosing the node that makes the most progress toward the ultimate destination of the message. For other protocols, this is extended to include other criteria. This option has the advantage of minimizing path dilation when no attacker is present.

Random selection is robust against a wide variety of attackers, since it does not give credence to the location information contained in the CTS. The larger the pool of forwarding candidates, the less likely that a neighboring attacker performing a CTS rushing attack or masquerading as a legitimate node is chosen. Performance suffers, however, since progress toward the destination is erratic. Compared with the impact of a black hole attack, this is most likely an acceptable tradeoff.

The originator of a message may choose more than one candidate to relay messages along multiple paths. This redundancy lessens the impact of attackers met along the way, though if a fixed number of attackers is present, the higher cost may be justified by its effectiveness.

4. **Omit Location** \in {yes, no} *Line 1*

Even when selecting among multiple responses in the collection window, an attacker can manipulate a choice if it is made by priority. Since the ORTS includes the ultimate destination, an attacker can fabricate an optimal location for inclusion in its CTS to maximize its chances of being selected.

An option to omit the source and destination locations in the ORTS message mitigates this threat. In this case, the neighbors of S cannot determine whether they are in the forwarding area, nor how close they lie to the line \overline{SD} . Therefore, all neighbors respond by setting timers proportional to their remaining energy only. The ORTS sender then chooses the relay according to the previous configuration setting.

When the DATA message is relayed to the selected node, it must contain the destination’s location to enable subsequent routing.

Omitting the destination does not eliminate the threat of a black hole attack, since an adversary may infer the ultimate destination from a stream of messages using traffic analysis. We do not consider that attack in this paper.

Note that during protocol operation, both a sender and its neighbors (forwarding candidates) retain some state. It is transient, however, since it need not be retained after the message is relayed. For this reason we classify SIGF-0 as stateless.

The configuration options presented give SIGF-0 robustness against a black hole caused by CTS rushing. They are similar enough to IGF to allow a smooth, runtime transition between option settings, according to the current attack situation. We are exploring the dynamic transition between settings, and between protocols in future work.

4.2 SIGF-1: Local-State Secure IGF

SIGF-1 builds on the capabilities and operation of SIGF-0, while aiming to further reduce the chances of selecting an attacker as the next-hop relay. By keeping some limited information about its current state and statistics of neighbor performance, a node can also defend against Sybil attacks. This state is summarized by a per-neighbor reputation value that influences the choice of forwarding candidates.

Since the state kept is not shared with neighbors, there is no overhead associated with initialization, synchronization, or repair. By limiting the information to that which can be verified locally, the protocol avoids state corruption attacks. Further, neighborhood dynamics due to mobility, failure, or transient communication are still considered.

We classify state kept in SIGF-1 in three categories: data about the local node, statistics about neighboring nodes, and values derived from both together. Each is presented below.

For the local node, we maintain T , the total number of messages sent by the node to all neighbors. It is used to calculate derived values for each neighbor. Nodes also have a small buffer B in which recently relayed messages are stored.

For each neighbor N among those discovered dynamically (i.e., **neighbor tables are not exchanged**), we keep the following:

1. N_{sent} = number of messages sent to neighbor N for forwarding. It is increased by one each time N is selected as the next-hop relay.
2. $N_{forward}$ = number of messages forwarded by neighbor N on this node’s behalf. This is counted by overhearing a message on its retransmission by node N to a downstream node, albeit imperfectly due to collisions and collusion.
3. $N_{location}$ = last claimed location of node N in its CTS message.
4. N_{delay} = average delay between relaying a message to node N and overhearing the subsequent relay of the same.

After transmitting a message to a neighboring node, a copy of the message is stored in the message buffer B , along with a timestamp. If the message is overheard on its relay to a downstream node, the difference between the recorded and

current times is calculated and the message is flushed from the buffer. $N_{forward}$ and N_{delay} are updated as described above. In case the buffer fills due to message loss or failure to overhear the relay, the oldest message is replaced and the associated N_{delay} is updated with a fixed maximum delay D .

From the data collected during routing, other values are derived which combine to determine a node’s reputation. These are also maintained per-neighbor as they are discovered.

5. $N_{success} = \frac{N_{forward}}{N_{sent}}$ = forwarding success ratio, a measure of reliability. Neighbors which always (verifiably) forward messages achieve high success ratios. When first discovered and until the neighbor forwards a message, it is given a neutral initial value of 0.5.
6. $N_{fairness} = \frac{T - N_{sent}}{T}$ = forwarding fairness ratio, a measure of the distribution of relay choices among neighbors. This promotes dispersion of next-hop relay choices among similarly performing neighbors, and reduces the likelihood of selecting an attacker even before its misbehavior is detected. Initial value is 0.5.
7. $N_{consistency}$ = a consistency score based on the variance of neighbor N ’s claimed locations. We assume legitimate nodes move little or rarely, so when $N_{location}$ changes, the score is decreased additively to penalize nodes which are either moving constantly or lying about their locations. When the claimed location remains the same, a small additive reward is granted, increasing the score. The consistency score saturates so as always to be in the interval $[0, 1]$. A neutral initial value of 0.5 is assigned.
8. $N_{performance} = \frac{D - N_{delay}}{D}$ = forwarding performance of the neighbor in terms of the maximum delay D , a static system parameter. This favors nodes which are able to quickly relay messages, due to light congestion and correct behavior, and penalizes nodes which are heavily congested or deliberately delay or drop messages. A neutral initial value of 0.5 is assigned.

Each neighbor is assigned a reputation R comprising a weighted linear combination of the above computed values:

$$R = \alpha N_{success} + \beta N_{fairness} + \gamma N_{consistency} + \zeta N_{performance} \quad (1)$$

The terms are weighted according to the network designer’s choice, with the limitation that all weights must sum to unity. All terms and the computed reputation are in the interval $[0, 1]$. The reputation is not shared externally; it is used only on the local node for ranking forwarding candidates.

SIGF-1 allows the weights to be assigned flexibly so the designer can favor some neighbor properties over others. The weights may also be adjusted dynamically based on current conditions. For example, a high weight for forwarding success ratio α may improve performance during a black hole attack by degrading attackers’ reputations more quickly. Table 1 summarizes the configuration parameters of SIGF-1.

SIGF-1 builds upon the stateless algorithm and protocol. All the options described above for SIGF-0 are still available. The Forwarding Area, Omit Location, and Collection Window settings are orthogonal, although a window for only one CTS responder does not provide any real choice of forwarding candidate. The key interaction is the use of reputation for choosing among eligible candidates.

Parameter	Description
α	Forwarding success weight
β	Forwarding fairness weight
γ	Location consistency weight
ζ	Forwarding performance weight
$R_{threshold}$	Reputation threshold

Table 1: System parameters for SIGF-1 to be determined statically by the network designer, or dynamically at runtime.

A reputation threshold $R_{threshold}$ is used to cull undesirable relays before applying the Forwarding Candidate Choice policy. All responders with reputations below the threshold are eliminated from consideration. Then the next-hop is chosen depending on the option in use: the *first* (i.e., earliest) responder, a *random* responder, or the responder with highest routing *priority* (based on distance to destination, etc.).

The threshold is a system parameter that allows nodes to avoid wasting energy sending a message to a neighbor with known poor performance, even if it claims to be the best route. The tradeoff is that a high threshold may cause premature routing failure by eliminating too many neighbors from consideration. For this reason, if no responders’ reputations are above $R_{threshold}$, we select the node with the highest reputation, even if it is a sub-optimal route. This favors routing success over performance when under attack.

Our experiments show that SIGF-1 effectively defends against the black hole, selective forwarding, and Sybil attacks when the next hop is selected using reputation. When an attacker drops messages, its reputation degrades quickly, as desired.

4.3 SIGF-2: Shared-State Secure IGF

Protection against attacks in SIGF-0 and SIGF-1 is gained by adding nondeterminism to the already dynamic forwarding candidate selection. However, some attacks still result in poor performance, since they go beyond the protections afforded by probabilistic, fully decentralized means.

SIGF-2 addresses this limitation by using state that is shared among neighbors for cryptographic operations. This provides guarantees for authenticity, confidentiality, integrity, and freshness that some other secure routing protocols provide (discussed in Section 6), but in the framework of a protocol family that can also operate without them. It builds upon SIGF-0 and SIGF-1, and inherits their configuration options.

The state required for use of SIGF-2 is described below along with the protocol configuration options. Many key pre-distribution or online key establishment protocols have been proposed, many of which are suitable for supporting this protocol. One example is LEAP [27], which provides both pairwise-shared keys between neighbors and neighborhood-shared keys for broadcast.

Each option is described below, including its shared-state requirements.

1. **Message Authentication** \in {all messages, only DATA, none}

Authenticating messages cryptographically ensures that they originate from a neighbor with which a node has

Protocol	General Approach	Corruption	Wormhole	HELLO flood	Black hole	Sybil	Replay DoS
IGF	Dynamic Binding	✓	✓	✓	–	–	–
SIGF-0	Nondeterminism	✓	✓	✓	✓	–	–
SIGF-1	Local Reputation	✓	✓	✓	✓	✓	–
SIGF-2	Cryptography	✓	✓	✓	✓	✓	✓

Table 2: Attacks resisted by IGF and SIGF protocols.

pre-shared information. This prevents an outside attacker from entering the network and being able to inject arbitrary messages. Using an appropriate key, a message authentication code (MAC) is computed over the header and payload and is appended to the message before transmission. Message integrity is provided by the same mechanism.

All protocol messages (ORTS, CTS, DATA, ACK) may be authenticated, or only the DATA portion. The latter has lower computation and communication overhead, but does not prevent replay attacks. It may, however, prevent an attacker from hijacking a protocol exchange to insert false data.

Note that message authentication does not prevent compromised nodes from participating in this or any other protocol. Since they possess all the security information of the original nodes, they may send any authenticated messages that the original nodes could.

CTS, DATA, and ACK authentication uses a shared key between the ORTS sender and the selected relay. When authenticating the ORTS message, a broadcast key must be used that is shared with all potential forwarding candidates in the neighborhood.

2. Message Sequencing $\in \{\text{yes, no}\}$

When message sequencing is enabled, protocol messages include a monotonically increasing sequence number s . A receiver accepts a message from neighbor N only if $s > N_{seq}$, the highest sequence number verifiably received from N . This ensures that each message is fresh and prevents an attacker from capturing and replaying old messages. It requires that N_{seq} be stored for each neighbor and updated upon each reception of an authentic message.

Message sequencing only provides defense against replay attacks if authentication is also in use. Otherwise, an attacker can simply change the sequence number when replaying a message, and it will not be detected at the receiver.

3. Payload Encryption $\in \{\text{yes, no}\}$

Payload encryption uses a key shared between an ORTS sender and the selected relay to conceal the contents of a DATA message from eavesdropping by attackers. This may also help to thwart traffic analysis based on semantic contents of messages.

The use of authentication and sequencing in SIGF-2 prevents message injection by outsiders, since they do not possess the keys to create valid MACs. Attackers also may not replay ORTS and CTS messages to cause denial of service

from spuriously invoking or aborting the protocol. Messages with old sequence numbers are dropped.

SIGF-2 does not by itself prevent compromised nodes from creating a black hole or other attack described in previous sections. By layering it atop SIGF-0 and SIGF-1 it retains these defenses.

End-to-end cryptographic protections may be employed at a higher level in the protocol stack. Such mechanisms would affect only the payload of the DATA message transparently to SIGF.

4.4 Discussion and Comparison

SIGF-0 inherits resistance to state corruption, wormhole, and HELLO flood attacks. In addition, it provides robustness against a black hole by CTS rushing. However, its effectiveness is reduced when an attacker creates multiple identities and responds to an ORTS with several CTS messages. Even with random selection from a large window, this greatly increases the chances of selecting the attacker.

SIGF-1 adds resistance to a Sybil attack by exploiting locally generated information about neighbors. The reputation calculation helps to distinguish between the stable, well-behaved legitimate neighbors and the attackers that lie about locations or do not forward packets reliably.

An attacker that replays others' messages can still mount a denial of service attack. To partially address this, we allow SIGF-2 to use state that is shared with its neighbors. Though the overhead is greater, this allows for cryptographic guarantees of authenticity, integrity, freshness, and confidentiality. It does not prevent flooding-type denial of service attacks, but mitigates against attackers using the protocol itself to cause disruption.

Even for authentic messages in SIGF-2, nodes do not completely trust neighbors. Methods of SIGF-0 and SIGF-1 for sampling among multiple candidates and ranking according to reputation are used to limit the impact of a compromised node.

Using overhearing to count the number of messages forwarded downstream ($N_{forward}$) in SIGF-1 is imperfect due to collisions, communication asymmetries, and collusion [19]. In the last case, an attacker may circumvent the check by relaying the message to a nonexistent node or to an accomplice, which drops it. The legitimate node's forwarding check is satisfied, and so the attack goes undetected. This more clever blackhole attack is detected by checking relay destinations against a cryptographically verified two-hop neighbor table. However, this additional state requires the use of SIGF-2.

Table 2 summarizes which attacks from Section 3 are addressed by the IGF and SIGF protocols.

The SIGF family of protocols is designed to provide several incremental steps between IGF and symmetric-cryptography-based routing protocols. This gives the network

designer flexibility to choose a protocol statically based on application security requirements and available resources. They can also be selected dynamically by control logic that remains for future work.

SIGF-0 reduces to emulating IGF operation when the forwarding area is a 60° sextant, the collection window is short enough to contain one response, the first response is selected, and the destination is included in the ORTS.

A node using SIGF-0 can dynamically change to SIGF-1 if notified by out-of-band means that attackers are present, or upon detecting degraded performance. Changing back again is as simple as releasing the state collected. SIGF-2 requires that keys be shared *a priori*, and so may not be available for dynamic selection at runtime if the network started completely statelessly. We investigate these issues in future work.

5. EVALUATION

To evaluate the performance of our secure routing protocol, we implemented SIGF in GloMoSim, a wireless simulator for sensor, ad hoc, and mobile networks. GloMoSim models the communication architecture from physical-layer bit transmissions, including signal interference and attenuation patterns, up the stack to application-layer traffic loads. Our system parameters are listed in Table 3.

Terrain	150 x 150 meters
Number of Nodes	196
Node Placement	Grid + $\mathcal{N}(0, 16)$ noise
Application	CBR streams
Payload Size	32 bytes
Simulation Length	100 packets, 10 runs
Radio Range	40 meters
Radio Bandwidth	200 kb/s

Table 3: GloMoSim simulation parameters.

For our experiments, we configured a terrain of 150 square meters, with 196 sensor nodes having communication radii of 40 meters. The terrain was subdivided uniformly into 196 cells. A node was placed at the center of each and then perturbed using a Gaussian distribution with standard deviation of four meters. We limited the duration of the CBR streams to 100 packets to emulate the type of traffic expected in low-bandwidth networks, and to avoid swamping initial reputation transients with steady-state behavior. Data shown in the graphs are the mean of ten simulation runs. Figure 5 shows the final node locations and labels for the source S , destination D , and attackers A1–4 used in the experiments.

We evaluated the performance of SIGF-0, SIGF-1, and SIGF-2 under six scenarios. In the base system test, we compare GF [5], DSR [11], IGF, and SIGF without any attacks. Then we evaluate SIGF under black hole, selective forwarding, and Sybil attacks. Denial of service attacks are considered in the last two experiments.

Configurations used for the experiments that follow are shown in Table 4. The labels given there are used in the discussions and legends of figures that follow.

5.1 Base System (No Attacks)

We consider the first set of experiments as a baseline for comparison. It tests many-to-many constant bit rate (CBR)




Figure 5: Locations for 196 nodes in GloMoSim 150x150m field. Nodes are first placed uniformly on a grid, then perturbed by a $\mathcal{N}(0, 16)$ distribution.

Label	Configuration details
SIGF-0 or SIGF-0-priority	60° forwarding area, fixed 5 ms collection window, choose by priority, include destination
SIGF-0-random	SIGF-0, but with random candidate selection
SIGF-1 or SIGF-1-reputation	SIGF-0 limited to high reputation neighbors, $\alpha = \frac{5}{8}$, $\beta = \frac{1}{8}$, $\gamma = \frac{1}{8}$, $\zeta = \frac{1}{8}$, $R_{threshold} = 0.45$
SIGF-1-random	SIGF-1, but with random candidate selection of nodes above $R_{threshold} = 0.45$
SIGF-2	SIGF-0 and SIGF-1 with all messages authenticated, message sequencing, and payload encryption

Table 4: Experimental protocol configurations.

traffic flows, which mimic the periodic point-to-point communication expected in such systems, for example from an event of interest back to a base station.

Results show that under increasing traffic load, SIGF modestly increases the overhead from message exchange and the end-to-end delay, but maintains high packet delivery ratios.

From Figure 6(a) we see that GF, IGF, and SIGF have comparable delivery ratios (90–100%) under light traffic load. When traffic flow rates increase to more than 7 packets/second per CBR flow, the network begins to suffer congestion in all protocols except IGF. Performance in GF degrades along limited intersecting routes, suffering additional congestion caused by neighbor table update beacons. SIGF suffers congestion since multiple CTS responses are collected by each ORTS sender. DSR has significant message loss from its flooded route discovery packets.

IGF saves in communication overhead because it does not require beaconing as in GF (see Figure 6(b)). Here the overhead packets are all MAC control packets including ORTS,

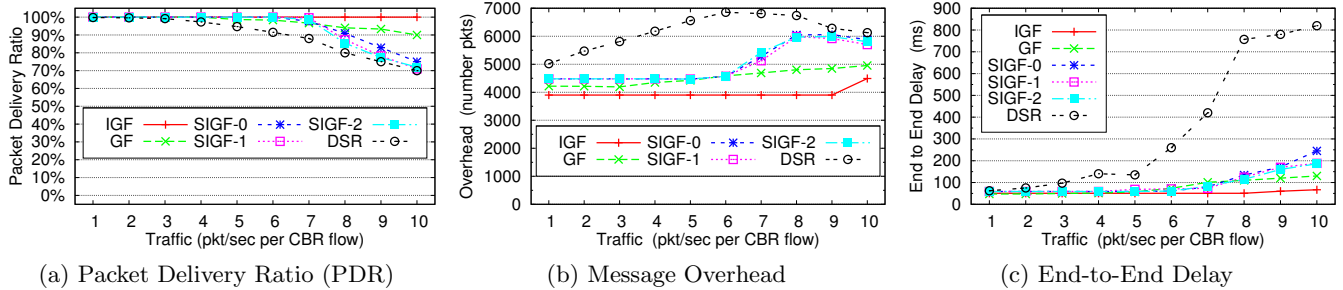


Figure 6: Baseline performance of routing protocols under increasing CBR traffic load, with no attacks.

CTS and ACK packets. Under light traffic loads, SIGF has similar communication overhead as GF, about 15% higher than IGF. As traffic loads increase, congestion increases the number of MAC layer collisions in IGF, GF and SIGF, resulting in retransmission attempts that add to the overhead. In particular, for SIGF the number of CTS packets increase quickly. DSR has more overhead because of route discovery packets. Its overhead ultimately diminishes because packet loss and the failure of route discovery packets to return to the source lead to fewer transmission attempts as messages are dropped early.

Local routing decisions introduce less end-to-end delay compared with routing protocols that require complete paths between a source and destination *a priori*. Figure 6(c) shows that IGF and SIGF have significantly lower end-to-end delay than DSR because DSR suffers latency awaiting the return of route discovery packets. This effect becomes less apparent in DSR under heavy traffic because DSR’s low delivery ratio leads to fewer packets contributing to this metric. Besides that, we also can see that SIGF causes only a gradual increase in end-to-end delay (from 59 to 188 ms for SIGF-1) after hitting congestion, even though it collects multiple CTS packets for every ORTS packet.

In summary, IGF has very good performance without attacks—better than GF. The SIGF protocols add minimal overhead, and though there is little to distinguish them in the baseline results, differences become clear as we add attacks in the next sections.

5.2 Black Hole Attack

To create a black hole, attackers rush their CTS responses (as described in Section 3.2.1) so that they are received first by the ORTS sender. If an attacker is selected as the next hop, it simply drops the packet (i.e., it does not transmit to a colluding or nonexistent node, which can only be detected by SIGF-2). We deal with incorrect locations in the Sybil attack experiments—here the locations reported are correct.

To eliminate the impact of network congestion, only a single CBR stream is considered. Node S sends a stream of packets to node D , shown in Figure 5. One-hop neighborhoods of S and D , and the direct line between them are also shown in the figure.

SIGF-0-random and SIGF-0-priority refer to the configurations shown in Table 4. SIGF-1-random and SIGF-1-reputation, also detailed in the table, discard responders whose reputations fall below $R_{threshold} = 0.45$. The former selects among the remaining nodes randomly, while the latter chooses the remaining node with the highest routing pri-

ority (based on distance, energy, etc). In both protocols, if no nodes have reputations that exceed the threshold, the node with the highest reputation is chosen.

We study the effect of the number of attackers and their locations on the packet delivery ratio in four scenarios. Figure 7 shows the performance of IGF and SIGF, grouped by the particular attack scenario. Results are nearly identical for CBR flows of 1–10 packets per second, hence the data shown are for six packets per second. Error bars in the graph show 95% confidence intervals on the mean.

In the first scenario, attacker $A2$ is near the routing path from source S to destination D (refer to Figure 5). However, it is not an optimal relay compared with other nodes, for example $A3$, which lies closer to the shortest geographic routing path.




Figure 7: Performance under black hole attack scenarios, six pkts/sec CBR flow from S to D .

As shown in Figure 7, under a black hole attack the Packet Delivery Ratio (PDR) of IGF becomes zero—it is unable to deliver a single packet, since the attacker is always the first responder. SIGF-0-priority, SIGF-1-reputation, and SIGF-2 all achieve better than 99% PDR: the former two protocols do not select attacker $A2$ since it is not an optimal choice, while the latter discards inauthentic messages. Using the random selection method in SIGF-0-random and SIGF-1-random degrades performance, since $A2$ will sometimes be chosen to relay the message. Still, the PDR is maintained at 50% and 79%, respectively.

In the second scenario, attacker $A3$ creates a black hole. From the node distribution, $A3$ is seen to lie near the optimal route from S to D . Hence, SIGF-0-priority performs very poorly, with zero PDR, since it always selects the at-

tacker as the next-hop relay. SIGF-1-reputation achieves a very high PDR of 98%, even though it also selects A3, since the reputation degrades quickly causing other nodes to be selected instead.

The combination of attackers A2 and A4 in the third scenario shows the cumulative effect of two black holes along the path from S to D . Since neither attacker is an optimal relay, SIGF protocols using priority, reputation, or authentication maintain PDR of 100%. Random selection in SIGF-0-random suffers most, since packet loss is incurred at two hops along the path. Performance in SIGF-1-random is unchanged at 73%, indicating that the attackers' reputations eventually degrade below $R_{threshold}$ and so cease to be among the neighbors chosen at random to relay.

Finally, attackers A1 and A3 in scenario four are both optimal relays. Performance degrades for the randomized and reputation-based protocols only slightly, since the attackers' reputations degrade to allow other nodes to be selected more frequently.

We note that unlike the other protocols, results for SIGF-2 assume an outsider is performing the black hole attack. In an attack by a compromised node, messages are authentic, and the protocol therefore performs as SIGF-1-reputation does—which is nearly as good.

In summary, SIGF protocols continue to deliver packets successfully when neighbors perform black hole attacks. Success rates vary depending on the amount of state and mechanisms used: SIGF-0 provides some defense with low PDRs (0–43%), SIGF-1 achieves moderate PDRs (70–99%), and SIGF-2 provides the best performance (100%).

5.3 Selective Forwarding Attack

If an attacker drops all messages completely, as in the black hole attack, it runs the risk that neighboring nodes can quickly conclude that an attack is under way and use other routes to avoid the attacker. It is more difficult to detect the attack if messages are selectively suppressed [13].

In this experiment, node A3 lies on the path of messages from S to D and mounts a selective forwarding attack. In Figure 8, successful packet delivery is plotted against an increasing packet drop ratio by the attacker. Error bars show 95% confidence intervals on the mean.

From zero to 100% dropped packets, IGF and SIGF-0-priority decline linearly from 100% to zero PDR. Randomized protocols SIGF-0-random and SIGF-1-random show a




Figure 8: Performance under selective forwarding attack by A3 for increasing packet drop ratios.

greater robustness, but still decline to 43% and 76%, respectively. The latter levels off due to its limited use of reputation. Delivery success for SIGF-1-reputation dips to 82% when the attacker drops 30% of packets, but improves thereafter since the packet loss is sufficient to degrade A3's reputation with its neighbors. Despite 50% dropped packets, SIGF-1-reputation has recovered to 96% PDR. SIGF-2 discards inauthentic messages, reliably achieving 100% PDR.

Here we clearly see the ability of SIGF-1-reputation to adapt to worsening attacks, using history to learn to avoid any unproductive neighbors. All the SIGF protocols react smoothly, without discontinuities or phase changes that may lead to unpredictable runtime behavior.

5.4 Sybil Attack

Now we evaluate our secure routing protocol under a Sybil attack by node A3. Figures 9–11 show the experimental results from the different scenarios we describe.

In the first scenario, attacker A3 creates six Sybil nodes randomly located about itself in a circle with a radius of the radio transmission range, the *Sybil distribution radius*. Each virtual node performs a black hole attack when the attacker receives an ORTS message. The locations of its virtual nodes are fixed to improve their reputations, since location inconsistency is penalized according to weight γ in Equation 1.




Figure 9: Performance under Sybil attack by A3, with six fixed-location virtual nodes.

Despite the Sybil black hole attack, SIGF-2 and SIGF-1-reputation achieve high packet delivery ratios, as shown in Figure 9. In SIGF-2, the attacker and its virtual Sybil nodes fail authentication, hence PDR is near 100%. In SIGF-1-reputation, Sybil nodes' reputations degrade quickly because they drop or modify the packets, resulting in a PDR of about 84%. Randomized protocols fare worse, but still achieve 26% and 35% PDRs. Overall, delivery ratios are less than in the single-node black hole attacks (Section 5.2) due both to more attackers and to network congestion caused by the Sybil neighbors.

When the number of virtual Sybil nodes increases, the delivery ratio is reduced because a Sybil node is more likely to be chosen as the next-hop relay. In the second scenario we simulated an increasing number of fake Sybil nodes to determine their impact on performance. Figure 10 shows the results. Although delivery ratios decline overall as the attacker uses more Sybil nodes, SIGF-1-reputation stabilizes for more than four Sybil nodes at about 80% PDR.




Figure 10: Performance under Sybil attack by A3, with increasing number of virtual nodes.

An attacker can maximize the impact of its virtual Sybil nodes by “locating” them entirely within the forwarding area of a nearby message stream, if possible. In the last scenario, we examine the impact on delivery ratio of increasing the Sybil distribution radii.

When the Sybil distribution radius is small, the attack is more effective if the attacker is close to the optimal forwarding path for a message stream. Such is the case here, in which A3 is the Sybil attacker. We observe a clustered PDR of about 30–34% for all but IGF and SIGF-2 (see Figure 11). Larger radii decrease the attack’s effectiveness (PDR improves to 58% or better) since fewer virtual Sybil nodes are in contention for relaying. These effects would be opposite for an attacker farther away from a message stream.




Figure 11: Performance under Sybil attack by A3, when virtual nodes are distributed in circles of increasing radii.

Together, these scenarios show that SIGF-0 and SIGF-1 can defend against Sybil attacks without requiring the initialization, synchronization, and state maintenance overhead of SIGF-2’s use of authentication. Although performance is better for SIGF-2, this may pose an acceptable tradeoff if the threat of Sybil attacks is low.

5.5 ORTS Replay DoS Attack

An attacker may capture and replay an ORTS message to cause a denial of service attack. For each ORTS message, this monopolizes the channel which may not be used except for collection of CTS messages from neighboring nodes.

In this experiment, node A3 replays an old ORTS message

every 100 ms while messages are in transit between S and D. Figure 12 shows that IGF, SIGF-0, and SIGF-1 are unable to defend against the attack, with less than 8% PDR in all cases. The congestion caused by the attacker’s denial of service causes almost all packets to be dropped in the neighborhood of the attacker.




Figure 12: Performance under ORTS Replay attack by A3, which replays an old ORTS every 100 ms.

Only SIGF-2 can determine that the message is inauthentic by examining the sequence number contained in it. The congestion causes a mere 10% loss of PDR. As with many denial of service attacks [24], defense against the ORTS Replay is difficult without the stronger guarantees of SIGF-2.

5.6 CTS Replay DoS Attack

In a CTS Replay attack, a node attempts to disrupt forwarding by causing other nodes to abort the protocol early. Or, the attacker may attempt to damage the reputation of a neighboring node by replaying old CTS messages, even if the neighbor is not currently responding.

In this experiment, for each ORTS message node A3 replays a legitimate CTS overheard from a neighbor. Figure 13 shows that this attack is much less damaging than the previous one. Although IGF and SIGF-0-priority are fooled, SIGF-2 and SIGF-1-reputation are not. In between are the protocols that select relays randomly. These suffer from the attack, but still allow 42% or 71% of messages to be delivered.

6. RELATED WORK

Although many secure routing protocols have been developed for ad-hoc networks, these are not directly applicable for several reasons. Some protocols (ARAN [23], SAODV [25], et al. [9, 17, 26]) use public-key cryptography, which is not memory and energy efficient enough for frequent use in sensor networks. Recent implementations of elliptic-curve algorithms on sensor devices may allow for their infrequent use, but symmetric cryptography-based algorithms are still desirable for their greater efficiency [18].

Some protocols use symmetric cryptography or hashing, but require maintenance of routing tables by online distance-vector algorithms (SEAD [7]) or on-demand multi-hop route discovery and caching (Ariadne [8], SRP [21]). For large-scale networks, this requires non-trivial consumption of memory and energy for the storage and update of routes to remote nodes. It also increases the “surface area” for security attacks.

Other work (e.g., SPINS [22], TinySec [12]) provides secure channels for use by otherwise unsecured protocols. They may be used to establish basic shortest-path routing trees

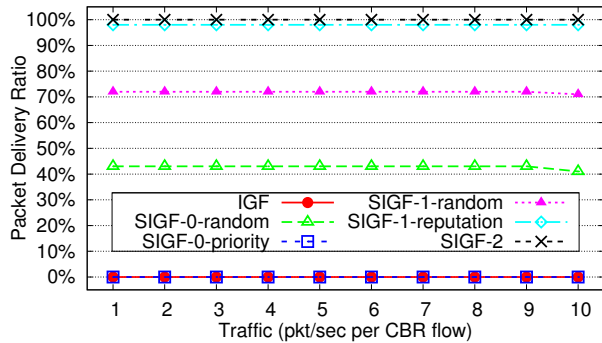


Figure 13: Performance during CTS Replay attack by node A3.

(as described in SPINS), but are inadequate defenses when nodes are compromised.

INSENS [3] is designed to tolerate node compromise and uses a variety of efficient mechanisms to establish routing. However, it is limited to routing upstream messages from nodes to base stations, using centralized topology collection and route computation.

Rather than maintain routing tables, SIGF chooses the next hop dynamically and non-deterministically. This contains the effect of compromise to a local neighborhood, increases robustness to node mobility and failure, and spreads energy drain more evenly across neighbors. Even without using symmetric cryptography, SIGF-0 and SIGF-1 achieve good performance under the attacks discussed in Section 3.

In addition to plain geographic forwarding (GF) [5] and IGF [1], on which SIGF is based, other geography-based routing algorithms have been proposed. GPSR and descendents [14, 15] extend GF to route around voids by traversing faces of a planar subgraph until greedy forwarding can resume. SIGF inherits a mechanism from IGF for handling forwarding failure, and many of the other techniques that have been proposed could be applied in the local or shared state contexts. ZRP [6] divides the network into variably-sized zones and allows different algorithms for intra- and inter-zone routing. These protocols are lightweight and efficient, but do not consider security.

7. CONCLUSION

We have presented SIGF (Secure Implicit Geographic Forwarding), a secure routing protocol family for wireless sensor networks that builds atop the inherently attack-containing, dynamic binding of IGF. Rather than maintain routing tables, SIGF chooses the next hop dynamically and non-deterministically. This contains the effect of compromise to a local neighborhood, increases robustness to node mobility and failure, and spreads energy drain more evenly across neighbors.

SIGF-0 keeps no state, but uses probabilistic means to avoid selecting an attacker for routing. SIGF-1 adds locally maintained reputations for dynamically discovered neighbors, using them to select well-behaved relays. SIGF-2 adds more traditional sequencing and cryptographic mechanisms for authentication, but at the greatest cost of resources.

We evaluated SIGF without attacks for base performance, and with black hole, selective forwarding, Sybil, and denial

of service attacks. We showed that even without using symmetric cryptography, SIGF-0 is able to defend against many attacks with no state, and SIGF-1 achieves high PDRs by maintaining reputations of neighbors. This allows efficient operation when no attacks are present, and good enough security when they are.

The basic IGF protocol handles voids and obstacles by shifting the sextant and retransmitting. Future studies will understand how this failure-recovery mechanism impacts the performance of the SIGF family of protocols. Practical systems need high density deployments, as evaluated in this paper, for reliability, power management, and sensing coverage. Future studies could evaluate lower densities to fully explore the behavior of SIGF.

8. ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their comments and Qiuhua Cao for her help on an early draft. This work was supported by NSF grants CNS-0615063 and CNS-0414870, and by ARO grant W9111NF-06-1-0204.

9. REFERENCES

- [1] B. Blum, T. He, S. Son, and J. Stankovic. IGF: A state-free robust communication protocol for wireless sensor networks. Technical Report CS-2003-11, Univ. of Virginia, Charlottesville, VA, 2003.
- [2] H. Chan, A. Perrig, and D. Song. Random key predistribution schemes for sensor networks. In *IEEE Symposium on Research in Security and Privacy*, pages 197–213, May 2003.
- [3] J. Deng, R. Han, and S. Mishra. A performance evaluation of intrusion-tolerant routing in wireless sensor networks. In *Proc. IEEE 2nd International Workshop on Information Processing in Sensor Networks (IPSN '03)*, Apr. 2003.
- [4] L. Eschenauer and V. D. Gligor. A key-management scheme for distributed sensor networks. In *9th ACM Conference on Computer and Communications Security*, pages 41–47, Nov. 2002.
- [5] G. G. Finn. Routing and addressing problems in large metropolitan-scale internetworks. Technical Report ISI/RR-87-180, ISI, Mar. 1987.
- [6] Z. J. Haas, M. R. Pearlman, and P. Samar. *The Zone Routing Protocol (ZRP) for ad hoc networks*. IETF MANET Internet Draft, July 2002.
- [7] Y.-C. Hu, D. B. Johnson, and A. Perrig. Secure efficient distance vector routing in mobile wireless ad hoc networks. In *Proceedings of the 4th IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, pages 3–13, June 2002.
- [8] Y.-C. Hu, A. Perrig, and D. B. Johnson. Ariadne: a secure on-demand routing protocol for ad hoc networks. In *Proceedings of the Eighth Annual International Conference on Mobile computing and Networking*, pages 12–23, Sept. 2002.
- [9] Y.-C. Hu, A. Perrig, and D. B. Johnson. Rushing attacks and defense in wireless ad hoc network routing protocols. In *ACM Workshop on Wireless Security (WiSe)*, pages 30–40, Sept. 2003.
- [10] IEEE Computer Society LAN MAN Standards Committee. *IEEE Std 802.11: Wireless LAN Medium*

- Access Control and Physical Layer Specifications*, Aug. 1999.
- [11] D. B. Johnson and D. A. Maltz. Dynamic source routing in ad hoc wireless networks. In Imielinski and Korth, editors, *Mobile Computing*, volume 353 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Mar. 1996.
- [12] C. Karlof, N. Sastry, and D. Wagner. TinySec: a link layer security architecture for wireless sensor networks. In *Proc. of the 2nd International Conference on Embedded Networked Sensor Systems (SenSys'04)*, pages 162–175, 2004.
- [13] C. Karlof and D. Wagner. Secure routing in wireless sensor networks: Attacks and countermeasures. In *First IEEE International Workshop on Sensor Network Protocols and Applications*, pages 1–15, May 2003.
- [14] B. Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (MOBICOM-00)*, pages 243–254, Aug. 2000.
- [15] Y.-J. Kim, R. Govindan, B. Karp, and S. Shenker. Geographic routing made practical. In *Proc. of the USENIX Symposium on Networked Systems Design and Implementation*, May 2005.
- [16] Y.-B. Ko and N. H. Vaidya. Location-aided routing (LAR) in mobile ad hoc networks. *Wireless Networks*, 6(4):307–321, 2000.
- [17] J. Kong, P. Zerfos, H. Luo, S. Lu, and L. Zhang. Providing robust and ubiquitous security support for mobile ad-hoc networks. In *9th International Conference on Network Protocols (ICNP'01)*, 2001.
- [18] D. Malan, M. Welsh, and M. Smith. A public-key infrastructure for key distribution in Tinyos based on elliptic curve cryptography. In *1st IEEE International Conference on Sensor and Ad Hoc Communications and Networks*, pages 71–80, Oct. 2004.
- [19] S. Marti, T. Giuli, K. Lai, and M. Baker. Mitigating routing misbehavior in mobile ad hoc networks. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (MOBICOM-00)*, pages 255–265, Aug. 2000.
- [20] J. Newsome, E. Shi, D. Song, and A. Perrig. The sybil attack in sensor networks: analysis & defenses. In *Proc. of 3rd IEEE/ACM Information Processing in Sensor Networks (IPSN'04)*, pages 259–268, Apr. 2004.
- [21] P. Papadimitratos and Z. Haas. Secure routing for mobile ad hoc networks. In *SCS Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS 2002)*, Jan. 2002.
- [22] A. Perrig, R. Szewczyk, V. Wen, D. Culler, and J. D. Tygar. SPINS: Security protocols for sensor networks. In *Proceedings of Seventh Annual International Conference on Mobile Computing and Networks MOBICOM 2001*, pages 189–199, July 2001.
- [23] K. Sanzgiri, B. Dahill, B. N. Levine, C. Shields, and E. M. Belding-Royer. A secure routing protocol for ad hoc networks. In *Proceedings of the IEEE International Conference on Network Protocols (ICNP)*, Nov. 2002.
- [24] A. D. Wood and J. A. Stankovic. Denial of service in sensor networks. *IEEE Computer*, 35(10):54–62, Oct. 2002.
- [25] M. G. Zapata and N. Asokan. Securing ad hoc routing protocols. In *Proc. ACM Workshop on Wireless Security (WiSe)*, pages 1–10. ACM Press, 2002.
- [26] L. Zhou and Z. J. Haas. Securing ad hoc networks. *IEEE Network*, 13(6):24–30, 1999.
- [27] S. Zhu, S. Setia, and S. Jajodia. LEAP: efficient security mechanisms for large-scale distributed sensor networks. In *Proceedings of the 10th ACM Conference on Computer and Communication Security*, pages 62–72, 2003.