

Chapter 21

Securing Mobile Ad Hoc Networks

Panagiotis Papadimitratos and Zygmunt J. Haas

Abstract

The vision of nomadic computing with its ubiquitous access has stimulated much interest in the mobile ad hoc networking (MANET) technology. These infrastructureless, self-organized networks, which either operate autonomously or as an extension to the wired networking infrastructure, are expected to support new MANET-based applications. However, the proliferation of this networking paradigm strongly depends on the availability of security provisions, among other factors. The absence of infrastructure, the nature of the envisioned applications, and the resource-constrained environment pose some new challenges in securing the protocols in ad hoc networking environments. The security requirements can differ significantly from those for infrastructure-based networks and the provision of security enhancements may take completely different directions as well. In this chapter, we study the schemes proposed to secure mobile ad hoc networks. We explain the primary goals of security enhancements, shed light on the commensurate challenges, survey the current literature on this topic, and finally introduce our approach to this multifaceted and intriguing topic.

21.1 Introduction

Mobile ad hoc networks comprise freely roaming wireless nodes that cooperatively make up for the absence of fixed infrastructure; that is, the nodes themselves support the network functionality. Nodes transiently associate with peers that are within the radio connectivity range of their transceivers and implicitly agree to assist in provision of the basic network services. These associations are dynamically created and torn down, often without prior notice or the consent of the communicating parties. MANET technology targets networks that can be rapidly deployed or formed in an arbitrary environment to enable communications or to serve a common objective dictated

by the supported application. Such networks can be highly heterogeneous, with various types of equipment, usage, transmission, and mobility patterns.

Secure communication, an important aspect of any networking environment, becomes an especially significant challenge in ad hoc networks. This is due to the particular characteristics of this new networking paradigm and due to the fact that traditional security mechanisms may be inapplicable.

The absence of a central authority deprives the network of the administrative and management services that would otherwise greatly facilitate its operation. MANET has to rely on continuous self-configuration, especially because of the highly dynamic nature of the network. Problems such as scheduling, address assignment, provision of naming services, and formation of network hierarchy cannot be solved by traditional centralized protocols. Instead, distributed operation is necessary in all aspects of network control, including basic security-related operations, such as the validation of node credentials. In the fully distributed and open environment of ad hoc networking, the provision of such services may not only incur a high overhead, but also provide additional opportunities for misbehaving nodes to harm the network operation.

In general, nodes participate in a protocol execution as peers, which implies that potentially any network node can abuse the protocol operation. As a result, it is fairly difficult to identify trustworthy and supportive nodes based on the network interaction. Additionally, determining the protocol or network components that have to be safeguarded is far from straightforward, something that makes the design of adequate security countermeasures even more difficult.

Meanwhile, the practically invisible (or nonexistent) administrative or domain boundaries make the enforcement of any security measures an even more complex problem. Migrating nodes may face varying rules even when they run the same application, as they move through different network areas and become associated with different groups of nodes. Or, they may lack the ground for the establishment of trust associations, that is, the establishment of some type of a secret, so that cryptographic mechanisms can be employed.

Below, we discuss in further detail the vulnerability of mobile ad hoc networks, clarify how security goals may have to be modified, and explain which types of solutions are plausible for different network instances. Although the discussion throughout a great part of the chapter applies to all types of ad hoc networks, it is important to realize that strictly not all solutions can be applied in all ad hoc networking environments. Moreover, it is necessary to emphasize the relative importance of addressing certain security issues, which can be considered, to some extent, as prerequisites

for solutions to other security problems. In the following sections, we present the challenges posed by the MANET environment, survey the relevant literature, identify the limitations of the proposed approaches, and suggest directions for future solutions.

21.2 Security Goals

The overall problem of securing a distributed system comprises the security of the networked environment and the security of each individual network node. The latter issue is important due to the pervasive nature of MANET, which does not allow us to assume that networked devices will always be under the continuous control of their owners. As a result, the physical security of the node becomes an important issue, leading to the requirement of tamper-resistant nodes [24], if comprehensive security is to be provided. However, security problems manifest themselves in a more emphatic manner in a networked environment and especially in mobile ad hoc networks. This is why in this work we focus on the network-related security issues.

Security encompasses a number of attributes that have to be addressed — availability, integrity, authentication, confidentiality, nonrepudiation, and authorization. These goals, which are not MANET-specific only, call for approaches that have to be adapted to the particular features of MANETs.

Availability ensures the survivability of network services despite misbehavior of network nodes; for instance, when nodes exhibit selfish behavior or when denial-of-service (DoS) attacks are mounted. DoS attacks can be launched at any layer of an ad hoc network. For example, an adversary could use jamming to interfere with communication at the physical layer or, at the network layer, it could disable the routing protocol operation, by disrupting the route discovery procedure. Moreover, an adversary could bring down high-level services. One such target is the key management service, an essential service for implementation of any security framework.

Integrity guarantees that an in-transit message is not altered. A message could be altered because of benign failures, such as radio propagation impairments, or because of malicious attacks on the network. Integrity viewed in the context of a specific connection, that is, the communication of two or more nodes, can provide the assurance that no messages are removed, replayed, reordered (if reordering would cause loss of information), or unlawfully inserted.

Authentication enables a node to ensure the identity of the peer node that it is communicating with. Without authentication, an adversary could masquerade a node, possibly gain unauthorized access to resources and sensitive information, and interfere with the operation of other nodes.

Confidentiality ensures that certain information is never disclosed to unauthorized entities. Confidentiality is required for the protection of sensitive information, such as strategic or tactical military information. However, confidentiality is not restricted to user information only; routing information may also need to remain confidential in certain cases. For example, routing information might be valuable for an enemy to identify and to locate targets on a battlefield.

Nonrepudiation ensures that the originator of a message cannot deny having sent the message. Nonrepudiation is useful for detection and isolation of compromised nodes. When node A receives an erroneous message from node B, A can use this message to accuse B and convince other nodes that B is compromised.

Finally, authorization establishes rules that define what each network node is or is not allowed to do. In many cases, it is required to determine which resources or information across the network a node can access. This requirement can be the result of the network organization or the supported application when, for instance, a group of nodes or a service provider wishes to regulate the interaction with the rest of the network. Another example could be when specific roles are attributed to nodes to facilitate network operation.

The security of mobile ad hoc networks has additional dimensions, such as privacy, correctness, reliability, and fault tolerance. In particular, the resilience to failures, which in our context can be the result of malicious acts, and the protection of the correct operation of the employed protocols are of critical importance and should be considered in conjunction with the security of the mobile ad hoc network.

21.3 Threats and Challenges

Mobile ad hoc networks are vulnerable to a wide range of active and passive attacks that can be launched relatively easily, because all communications take place over the wireless medium. In particular, wireless communication facilitates eavesdropping, especially because continuous monitoring of the shared medium, referred to as promiscuous mode, is required by many MANET protocols. Impersonation is another attack that becomes more feasible in the wireless environment. Physical access to the network is gained simply by transmitting with adequate power to reach one or more nodes in proximity, which may have no means to distinguish the transmission of an adversary from that of a legitimate source. Finally, wireless transmissions can be intercepted. An adversary with sufficient transmission power and knowledge of the physical and medium access control layer mechanisms can obstruct its neighbors from gaining access to the wireless medium.

Assisted by these opportunities that wireless communication offers, malicious nodes can meaningfully alter, discard, forge, inject, and replay control and data traffic; generate floods of spurious messages; and, in general, avoid complying with the employed protocols. The impact of such malicious behavior can be severe, especially because the cooperation of all network nodes provides for the functionality of the absent fixed infrastructure. In particular, as part of the normal operation of the network, nodes are transiently associated with a dynamically changing, over time, subset of their peers; that is, the nodes within the range of their transceivers or the ones that provide routing information and implicitly agree to relay their data packets. As a result, a malicious node can obstruct the communications of potentially any node in the network, exactly because it is entitled or even expected to assist in the network operation.

In addition, freely roaming nodes join and leave MANET subdomains independently, possibly frequently, and without notice, making it difficult in most cases to have a clear picture of the ad hoc network membership. In other words, there may be no ground for an *a priori* classification of a subset of nodes as trusted to support the network functionality. Trust may only be developed over time, although trust relationships among nodes may also change, when, for example, nodes in an ad hoc network dynamically become affiliated with administrative domains. This is in contrast to other mobile networking paradigms, such as Mobile IP or cellular telephony, where nodes continue to belong to their administrative domain in spite of mobility. Consequently, security solutions with static configuration would not suffice, and the assumption that all nodes can be bootstrapped with the credentials of all other nodes would be unrealistic for a wide range of MANET instances.

From a slightly different point of view, it becomes apparent that nodes cannot be easily classified as internal or external (i.e., nodes that belong to the network or not): nodes that are expected to participate and be dedicated to supporting a certain network operation and those that are not. In other words, the absence of an infrastructure impedes the usual practice of establishing a line of defense, separating nodes into trusted and not trusted. As a result, attacks cannot be classified as internal or external either, especially at the network layer. Of course, such a distinction could be made at the application layer, where access to a service or participation in its collaborative support may be allowed only to authorized nodes. In the latter example, an attack from a compromised node within the group (i.e., a group node under the control of an adversary) would be considered an internal one.

The absence of a central entity makes the detection of attacks a difficult problem, because highly dynamic large networks cannot be easily monitored. Benign failures, such as transmission impairments, path breakages,

and dropped packets, are a fairly common occurrence in mobile ad hoc networks and, consequently, malicious failures will be more difficult to distinguish. This will be especially true for adversaries that vary their attack pattern and misbehave intermittently against a set of their peers that also changes over time. As a result, short-lived observations will not allow detection of adversaries. Moreover, abnormal situations may occur frequently because nodes behave in a selfish manner and do not always assist the network functionality. It is noteworthy that such behavior may not be malicious, but only necessary when, for example, a node shuts its transceiver down to preserve its battery.

Most of the currently considered MANET protocols [30] were not originally designed to deal with malicious behavior or other security threats. Thus, they are easy to abuse. Incorrect routing information can be injected by malicious nodes that respond with or advertise nonexistent or stale routes and links. In addition, compromised routes (i.e., routes that are not free of malicious nodes), may be repeatedly chosen with the encouragement provided by the malicious nodes themselves. The result is that pairs of communicating end nodes will experience DoS and may have to rely on cycles of time-out and new route discovery to find operational routes, with successive query broadcasts imposing additional overhead. Or even worse, the end nodes may be easily deceived for some period of time that the data flow is uninterrupted, although no actual communication takes place. For example, an adversary may drop a route error message, hiding a route breakage, or forge network and transport layer acknowledgments.

Finally, mobile or nomadic hosts have limited computational capabilities, due to constraints stemming from the nature of the envisioned MANET applications. Expensive cryptographic operations, especially if they have to be performed for each packet and over each link of the traversed path, make such schemes implausible for the vast majority of mobile devices. Cryptographic algorithms may require significant computation delays, which in some cases would range from one to several seconds for low-end devices [5, 11]. These delays, imposed, for example, by the generation or verification of a single digital signature, affect the data rate of secure communication. More important, mobile devices could become targets of DoS attacks due to their limited computational resources. An adversary could generate bogus packets, forcing the device to consume a substantial portion of its resources. Even worse, a malicious node with valid credentials could frequently generate control traffic, such as route queries, at a high rate not only to consume bandwidth, but also to impose cumbersome cryptographic operations on a sizable portion of the network nodes.

21.4 Trust Management

The use of cryptographic techniques is necessary for the provision of any type of security services; mobile ad hoc networks are not an exception to this rule. The definition and the mechanisms for security policies, credentials, and trust relationships (i.e., the components of what is collectively identified as trust management) are a prerequisite for any security scheme. A large number of solutions have been presented in the literature for distributed systems, but they cannot be readily transplanted into the MANET context, because they rely on the existence of a network hierarchy and a central entity. Envisioned applications for the ad hoc networking environment may require a completely different notion of establishing a trust relationship, although the network operation may impose additional obstacles to the effective implementation of such solutions.

For small-scale networks, of the size of a personal or home network, trust can be established in a truly ad hoc manner, because relationships can be static and sporadically reconfigured manually. In such an environment, the owner of a number of devices or appliances can imprint them, that is, distribute their credentials along with a set of rules that determine the allowed interaction with and between devices [24]. The proposed security policy follows a master–slave model, with the master device being responsible for reconfiguring slave devices, issuing commands, or retrieving data. The return to the initial state can be performed only by the master device or by some trusted key escrow service.

This model naturally lends itself to represent personal area networking, in particular network instances such as Bluetooth® wireless technology [4], in the sense that within a piconet, the interactions between nodes can be determined by the security policy. The model can be extended by allowing partial control or access rights to be delegated, so that the secure interaction of devices becomes more flexible [25]. However, if the control over a node can be delegated, the new master should be prevented from eradicating prior associations and assuming full control of the node.

A more flexible configuration, independent of initial bindings, can be useful when a group of people wish to form a collaborative computing environment [9]. In such a scenario, the problem of establishing a trust relationship can be solved by a secure key agreement, so that any two or more devices are able to communicate securely. The mutual trust among users allows them to share or establish a password using an offline secure channel or perform a preauthentication step through a localized channel [1]. Then, they can execute a password-based authenticated key exchange over the nonsecure wireless medium. Schemes that derive a shared symmetric key could use a two- or a multiparty version of the password authenticated Diffie–Hellman key-exchange algorithm [3].

Human judgment and intervention can greatly facilitate the establishment of spontaneous connectivity among devices. Users can select a shared password or manually configure the security bindings between devices, as seen above. Furthermore, they could assess subjectively the security of their physical and networking environment and then proceed accordingly. However, human assistance may be impossible for the envisioned MANET environment with nodes acting as mobile routers, even though the distinction between an end device and a router may be only logical, with nodes assuming both roles. Frequently, the sole requirement for two transiently associated devices will be to mutually assist each other in the provision of basic networking services, such as route discovery and data forwarding. This could be so because mobile nodes do not necessarily pursue collectively a common goal. As a result, the users of the devices may have no means to establish a trust relationship in the absence of a prior context.

However, there is no reason to believe that a more general trust model would not be required in the MANET context. For instance, a node joining a domain may have to present its credentials to access an available service and, at the same time, authenticate the service itself. Similarly, two network nodes may wish to employ a secure mode of multi-hop communication and verify each other's identity. Clearly, support for such types of secure interaction, either at the network or at the application layer, will be needed.

A public key cryptosystem can be a solution, with each node bound to a pair of keys, one publicly known and one private. However, the deployment of a public key infrastructure (PKI) requires the existence of a certification authority (CA), a trusted third party responsible for certifying the binding between nodes and public keys. The use of a single point of service for key management can be a problem in the MANET context, especially because such a service should always remain available. It is possible that network partitions or congested links close to the CA server, although they may be transient, could cause significant delays in getting a response. Moreover, in the presence of adversaries, access to the CA may be obstructed or the resources of the CA node may be exhausted by a DoS attack. One approach is not to rely on a CA and thus abolish all the advantages of such a facility. Another approach is to institute the CA in a way that answers the particular challenges of the MANET environment.

The former approach can be based on the bootstrapping of all network nodes with the credentials of every other node. However, such an assumption would dramatically narrow the scope of ad hoc networking, because it can be applied only to short-lived mission-oriented, and thus, closed networks. An additional limitation may stem from the need to ensure a sufficient level of security, which implies that certificates should be refreshed from time to time, requiring, again, the presence of a CA.




Figure 21.1. Configuration of a Key Management Service Comprising n Servers

The service, as a whole, has a public/private key pair K/k . The public key K is known to all nodes in the network, whereas the private key k is divided into n shares s_1, s_2, \dots, s_n , with one share for each server. Moreover, each server has a public/private key pair K_i/k_i and knows the public keys of all nodes.

(Source: L. Zhou and Z.J. Haas, Securing Ad Hoc Networks, *IEEE Network Magazine*, November/December 1999. © 1999 IEEE) Used with permission.)

Alternatively, it has been suggested that users certify the public keys of other users. One such scheme proposes that any group of K nodes may provide a certificate to a requesting node. Such a node broadcasts the request to its one-hop neighborhood, each neighbor provides a partial certificate, and if sufficient K such certificates are collected, the node acquires the complete certificate [14, 29]. Another scheme proposes that each node select a number of certificates to store, so that, when a node wants the public key of one of its peers, the two certificate repositories are merged, and if a chain of certificates is discovered, the public key is obtained [13].

The solution of a key management facility that meets the requirement of the MANET environment has been proposed in [29]. To do so, the proposed instantiation of the public key infrastructure provides increased availability and fault tolerance. The distributed instantiation of the CA is equipped with a private/public key pair. All network nodes know the public key of the CA and trust all certificates signed by the CA's private key. Nodes that wish to establish secure communication with a destination query the CA and retrieve the required certificate, thus being able to authenticate the other end and establish a secret shared key for improved efficiency. Similarly, nodes can request an update from the CA (i.e., change their own public key and acquire a certificate for the new key).

The distributed CA is instantiated by a set of nodes (servers), as shown in Figure 21.1, for enhanced availability. However, this is not done through naïve replication, which would increase the vulnerability of the system, because the compromise of a single replica would be sufficient for the adversary to control the CA. Instead, the trust is distributed among a set of nodes that share the key management responsibility. In particular, each of the n servers has its own pair of public/private keys and they collectively share the ability to sign certificates. This is achieved with the use of threshold cryptography,

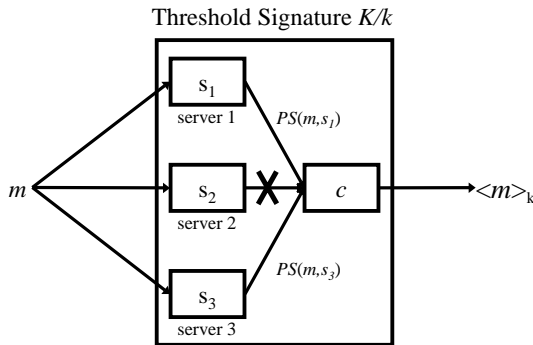


Figure 21.2. Calculation of a Threshold Signature

As an example, the service consists of three servers and uses a (3, 2) threshold cryptography scheme. K/k is the public/private key pair of the service and each server has a share s_i of the private key. To calculate the threshold signature on a message m , each server generates a partial signature $PS(m, s_i)$ and correct server 1 and server 3 forward their partial signatures to a combiner c . Even though server 2 fails to submit a partial signature, c is able to generate the signature $\langle m \rangle_k$ (m signed by the service private key k).

(Source: L. Zhou and Z.J. Haas, Securing Ad Hoc Networks, *IEEE Network Magazine*, November/December 1999. © 1999 IEEE) Used with permission.)

which allows any $t + 1$ out of n parties to perform a cryptographic operation, although t parties cannot do so. To accomplish this, the private key of the service, as a whole, is divided into n shares, with each of the servers holding one share. When a signature has to be computed, each server uses its share and generates a partial signature. All partial signatures are submitted to a combiner, a server with the special role to generate the certificate signature out of the collected partial signatures, as shown in Figure 21.2. This is possible only with at least $t + 1$ valid partial signatures.

The application of threshold cryptography provides protection from compromised servers, because more than t servers have to be compromised before it assumes control of the service. If fewer than $t + 1$ servers are under the control of an adversary, the operation of the CA can continue, because purposefully invalid partial signatures, contributed by rogue servers, will be detected. Moreover, the service provides the assurance that the adversary will not be able to compromise enough servers over a long period of time. This is done with the help of share refreshing, a technique that allows the servers to calculate new shares from the old ones without disclosing the private key of the service. The new shares are independent from the older ones and cannot be combined with the old shares in an attempt to recover the private key of the CA. As a result, to compromise the system, all $t + 1$ shares have to be compromised within one refresh period, which can be chosen appropriately short in order to

decrease vulnerability. The vulnerability can be decreased even further, when a quorum of correct servers detects compromised or unavailable servers and reconfigures the service. In other words, the quorum generates and distributes a new set of n shares, $t + 1$ of which need be combined now to calculate a valid signature. It is noteworthy that the public/private key pair of the service is not affected by share refreshing and reconfiguration operations, which are transparent to all clients.

The threshold cryptography key management scheme can be adapted further by selecting different configurations of the key management service for different network instances. For example, the numbers of servers can be selected according to the size or the rate of membership changes of the network; for a large number of nodes within a large coverage area, the number of servers should also be large, so that the responsiveness of the service can be high. Nodes will tend to interact with the closest server, which can be only a few hops away, or with the server that responds with the least delay. Another possibility is to alternate among the servers within easy reach of the client, something that can happen naturally in a dynamically changing topology. This way, the load from queries and updates will be balanced among different servers and the chances of congestion near one of the servers will be reduced. At the same time, the storage requirements can be traded off for interserver communication, by storing at each server a fraction of the entire database.

Additionally, the efficient operation of the CA can be enhanced when it is combined with secure route discovery and data forwarding protocols. Such protocols could, in fact, approximate the assumption of reliable links between servers in [29] even in the presence of adversaries. In particular, two of the protocols that will be discussed below — Secure Routing Protocol (SRP) and Secure Message Transmission (SMT) — lend themselves naturally to this model. Any two servers can discover and maintain routes to each other and forward service-related traffic, regardless of whether or not intermediate nodes are trusted.

21.5 Secure Routing

The secure operation of the MANET routing protocol is of central importance because of the absence of a fixed infrastructure. Instead, nodes are transiently associated and will cooperate with virtually any node, including those that could potentially disrupt the route discovery and data forwarding operations. In particular, the disruption of the route discovery may be an effective means to systematically obstruct the flow of data. Adversaries can respond with stale or corrupted route replies or broadcast forged control packets to obstruct the propagation of legitimate queries and routing updates.

However, the usual practice for securing Internet routing protocols [19] cannot be applied in the MANET context. The schemes proposed to secure Internet routing rely mainly on the existence of a line of defense, separating the fixed routing infrastructure from all other network entities. This is achieved by distributing a set of public keys/certificates, which signify the authority of the router to act within the limits of the employed protocol (e.g., advertise certain routes) and allow all routing data exchanges to be authenticated, not repudiated, and protected from tampering. However, such approaches cannot combat a malicious router disseminating incorrect topological information. More importantly, they are not applicable in the MANET context because of impediments such as the absence of a fixed infrastructure and a central entity.

Although the appropriate design could provide increased assurances of the availability of an online CA, the requirement to authenticate all nodes assisting the route discovery and data communication may not be practical. The interaction with the CA could become a limiting factor, because a node would have to acquire and validate the credentials of all intermediate network nodes along the discovered routes. Clearly, at least one route to the server has to be discovered before the node can contact the node instituting the CA server. But the problem is that, in the presence of adversaries, forged replies would still require the server's response to be validated. Another important limitation arises from the frequently changing topology and network membership, which would incur frequent queries addressed to the CA. In addition, congested links close to the server, although they may be transient or intermittent, could result in significant delays or even total failure to provide the certification services.

The protection of the route discovery process has been regarded as an additional Quality-of-Service (QoS) issue [28], by choosing routes that satisfy certain quantifiable security criteria. In particular, nodes in a MANET subnet are classified into different trust and privilege levels. A node initiating a route discovery sets the sought security for the route, that is, the required minimum trust level for nodes participating in the query/reply propagation. Nodes at each trust level share symmetric encryption and decryption keys. Intermediate nodes of different levels that cannot determine whether the required QoS parameter can be satisfied or decrypt in-transit routing packets drop them. This scheme provides protection (e.g., integrity) of the routing protocol traffic against adversaries outside a specific trust level.

An extension of the Ad hoc On-Demand Distance Vector (AODV) [20] routing protocol has been proposed [10] to protect the routing protocol messages. The Secure-AODV scheme assumes that each node has the certified

public keys of all network nodes, because intermediate nodes validate all in-transit routing packets. The basic idea is that the originator of a control message appends a RSA signature [23] and the last element of a hash chain [15] (i.e., the result of n consecutive hash calculations of a random number). As the message traverses the network, intermediate nodes cryptographically validate the signature and the hash value, generate the k -th element of the hash chain, with k being the number of traversed hops, and place it in the packet. Route replies are provided either by the destination or by intermediate nodes that have an active route to the sought destination.

A second proposal to secure AODV makes use of public key cryptography as well and operates in two stages — an end-to-end authentication and an optional secure shortest path discovery [7]. First, a signed route request propagates to the sought destination, which returns a signed response to the querying node. At each hop, for either direction, the receiving node validates the received control packet and forwards it after signing it. At the second stage, a shortest path confirmation packet is sent toward the destination, while now intermediate nodes sign the message in an onion-like manner to disallow changes of the path length.

21.5.1 The Secure Routing Protocol

The SRP [17] for mobile ad hoc networks provides correct end-to-end routing information over an unknown, frequently changing network, in the presence of malicious nodes. It is assumed that any two nodes that wish to employ SRP have a security association (SA), such as a symmetric shared secret key. Communication takes place over a broadcast medium and it is assumed that malicious nodes, which may concurrently corrupt the route discovery, cannot collude during a single route discovery. Moreover, we assume that nodes have a single data link interface, with a one-to-one correspondence between data link and IP addresses. Under these assumptions, the protocol is proven robust.

SRP provides one or more route replies, the SRP correctness of which is verified by the route geometry itself, and compromised and invalid routing information is discarded. The route request packets verifiably propagate to the destination, and route replies are returned to the querying node strictly over the reversed route, as accumulated in the route request packet. To guarantee this crucially important functionality, SRP employs an explicit interaction with the network layer (i.e., the IP-related functionality). Moreover, a number of novel features allow SRP to safeguard the route discovery operation, as explained below.

21.5.1.1 The Neighbor Lookup Protocol. The Neighbor Lookup Protocol (NLP) is an integral part of SRP responsible for the following tasks:

1. It maintains a mapping of Media Access Control and IP layer addresses of the node's neighbors.
2. It identifies potential discrepancies, such as the use of multiple IP addresses by a single data-link interface.
3. It measures the rates at which control packets are received from each neighbor, by differentiating the traffic primarily based on Media Access Control addresses.

The measured rates of incoming control packets are provided to the routing protocol as well. This way control traffic originating from nodes that selfishly or maliciously attempt to overload the network can be discarded.

Basically, NLP extracts and retains the 48-bit hardware source address for each received (overheard) frame along with the encapsulated IP address. This requires a simple modification of the device driver [27], so that the data link address is passed up to the routing protocol with each packet. With nodes operating in promiscuous mode, the extraction of such pairs of addresses from all overheard packets leads to a reduction in the use of the neighbor discovery and query/reply mechanisms for medium access control address resolution. Each node updates its neighbor table by retaining both addresses. The mappings between data-link and network interface addresses are retained in a table as long as transmissions from the corresponding neighboring nodes are overheard; a time-out period is associated with each entry removed from the table upon expiration.

NLP issues a notification to SRP in the event that according to the content of a received packet:

1. A neighbor used an IP address different from the address currently recorded in the neighbor table.
2. Two neighbors used the same IP address (i.e., a packet appears to originate from a node that may have spoofed an IP address).
3. A node uses the same medium access control address as the detecting node (in that case, the data link address may be spoofed).

Upon receiving the notification, the routing protocol discards the packet bearing the address that violated the aforementioned policies.

Even without cryptographic validation, NLP thwarts adversaries from presenting themselves at the routing layer as more than one node. This would have been possible if different IP addresses were inserted in or used as the source address of the control traffic the adversary relays or originates. However, the effectiveness of NLP relies on the fact that medium access control addresses are either hardwired or may be changed only with substantial latency. In the former case, NLP can provide strong assurances; in the latter one, it will be a significant line of defense, deterring, for example a malicious node from flooding the network with spurious traffic. Moreover, NLP can use cryptography and establish pair-wise security

associations (e.g., symmetric secret keys) among neighbors, so that nodes can authenticate traffic received from their neighbors. In any case, we should note that it is not of interest for SRP whether a relay node indeed presented itself with its actual IP address, but whether the node participated in the discovery of the route.

21.5.1.2 The Basic Secure Route Discovery Procedure. The querying node maintains a query sequence number, Q_{seq} , for each destination it securely communicates with. The monotonically increasing sequence number allows the destination to detect outdated route requests. At the same time, route requests are assigned a pseudorandom query identifier, Q_{ID} , which is used by intermediate nodes. Q_{ID} is statistically indistinguishable from a random number and thus unpredictable by an adversary with limited computational power. As a result, broadcasted fabricated requests will fail to cause subsequent legitimate queries to be dropped as previously seen, if, for example, the forged packets carry a higher sequence number.

Both Q_{ID} and Q_{seq} are placed in the SRP header, along with a message authentication code (MAC) that covers the shared key, $K_{S,T}$, and the protocol header. Fields that are updated as the packet propagates toward the destination, such as the accumulated addresses, are excluded from the MAC calculation.

Nodes compare the last entry in the accumulated route to the IP datagram source address, which belongs to the neighboring node that relayed the request. If there is a mismatch, or NLP provides a notification that the relaying neighbor violated one of the enforced policies, the query is dropped. Otherwise the Q_{ID} and the source and destination addresses are placed in the query table, so that previously seen queries are discarded. Fresh route requests are rebroadcast, with intermediate nodes inserting their IP address in the request packet.

The destination validates the integrity and freshness of queries originating from nodes it is securely associated with. It generates a number of replies that does not exceed the number of its neighbors, so that a malicious neighbor does not control more than one route. The reversed accumulated route serves as the source route of the reply packet, which is identified by Q_{ID} and Q_{seq} . The appended MAC covers the SRP header, including the source route. This way the source can be provided with evidence that the request had reached the destination and, in conjunction with the source route, that the reply was indeed returned along the reverse of the discovered route.

As the reply propagates along the reverse route, each intermediate node simply checks whether the source address of the route reply datagram is the same as the one of its downstream node, as determined by the route reply; if not, the reply is discarded. Ultimately, the source validates the

reply by first checking whether it corresponds to a pending query. Then, it is sufficient to validate the MAC, because the IP source-route already provides the (reversed) route itself.

21.5.1.3 Priority-Based Query Handling. To guarantee the responsiveness of the routing protocol, nodes maintain a priority ranking of their neighbors according to the rate of queries observed by NLP. The highest priority is assigned to the nodes generating (or relaying) requests with the lowest rate and vice versa. Quanta are allocated proportionally to the priorities and low-priority queries that are not serviced are eventually discarded. Within each class, queries are serviced in a round robin manner. Selfish or malicious nodes that broadcast requests at a very high rate are throttled back, first by their immediate neighbors and then by nodes farther from the source of potential misbehavior. Nonmalicious queries (i.e., queries originating from benign nodes that regulate in a nonselfish manner the rate of query generation) will be affected only for a period equal to the time it takes to update the priority (weight) assigned to a misbehaving neighbor. In the mean time, the round robin servicing of requests provides the assurance that benign requests will be relayed even amid a storm of malicious or extraneous requests.

21.5.1.4 The Route Maintenance Procedure. The route-error packets are source-routed to the end node along the prefix of the route that is being reported as broken. The intermediate upstream nodes, with respect to the point of breakage, check whether the source address of the route error datagram is the same as the one of their downstream node as reported in the broken route. Then, if there is no notification from NLP that the relaying neighbor violated one of the enforced policies, they relay the packet toward the source. In this case, NLP prevents an adversary that does not belong to the route, but lies at a one-hop distance from it, from generating an error message, because an inconsistency with the addresses already used (during the route discovery) by the actual downstream neighbor will be detected.

The notified source compares the source-route of the error message to the prefix of the corresponding active route. This way, it verifies that the provided route error message refers to the actual route and that it is not generated by a node that is not part of the route. The correctness of the feedback (i.e., whether it reports an actual failure to forward a packet) cannot be verified, though. As a result, a malicious node lying on a route can mislead the source by corrupting error messages generated by another node or by masking a dropped packet as a link failure. However, it can harm only the route it belongs to, something that was possible in the first place if it simply dropped or corrupted the data packets. To ensure that a non-operational route is detected or to verify that a route remains operational,

the source can utilize feedback from the trusted destination. This functionality can be provided by secure data communication protocols, such as the Secure Message Transmission (SMT) protocol. We discuss this in Section 21.6.1, with route maintenance providing complementary information.

21.5.1.4 The SRP Extension. The basic operation of SRP can be extended to allow for nodes other than the destination to provide route replies. This would be possible only under additional trust assumptions, when, for example, nodes sharing a common objective belong to the same group and mutually trust all the group members. In particular, this could be the case when all group members share a secret key.

Under this assumption, a querying node appends to each query an additional MAC calculated with the group key, which we call Intermediate Node Reply Token (INRT). The functionality of SRP remains as described above, with the following addition: each group member maintains the latest Q_{ID} seen from each of its peers and can thus validate both the freshness and origin authenticity of queries generated from other group nodes.

If a node other than the sought destination receives such a valid query, it can respond to the request if it has knowledge of a route to the destination in question. However, the correctness of such a route is conditional upon the correctness of the second portion of the route, which is provided by the intermediate node.

This functionality can be provided independently from and in parallel with the one relying solely on the end-to-end security associations. For example, it could be useful for frequent intragroup communication; any two members can benefit from the assistance of their trusted peers, which may already have useful routes.

21.6 Secure Data Forwarding

The frequent interaction with a CA and the frequent use of computationally expensive cryptographic tools are restrictive assumptions, especially for secure data-forwarding schemes. Such protocols must also take into account the inherent limitations of the MANET paradigm, exploit its features, and incorporate widely accepted and evaluated techniques to be efficient and effective. The above SRP for mobile ad hoc networks satisfies the above-stated goals.

However, SRP or any other underlying routing protocol cannot guarantee that the nodes along a correctly discovered route will, indeed, relay the data as expected. An adversary may misbehave in an intermittent manner (i.e., provide correct routing information during the route discovery stage and later forge or corrupt data packets during the data-forwarding stage). This is

exactly the function that is required by any secure data-forwarding protocol: to secure the flow of data traffic in the presence of malicious nodes, after the routes between the source and the destination have been discovered.

One of the solutions targeting the MANET environment proposes two mechanisms that detect misbehaving nodes and report such events, by maintaining a set of metrics reflecting the past behavior of other nodes [16].

To alleviate the detrimental effects of packet dropping, nodes choose the best route, comprised of relatively well-behaved nodes (i.e., nodes that do not have a history of avoiding forwarding packets). Among the assumptions in the above-mentioned work are a shared medium, bidirectional links, use of source routing (i.e., packets carry the entire route that becomes known to all intermediate nodes), and no colluding malicious nodes. Nodes operating in promiscuous mode overhear the transmissions of their successors and may verify whether the packet was forwarded intact to the downstream node. Upon detection of a misbehaving node, a report is generated, and nodes update the rating of the reported misbehaving node. The ratings of nodes along a well-behaved route are periodically incremented, although reception of a misbehavior alert dramatically decreases the nodes rating. When a new route is required, the source node calculates a path metric equal to the average of the ratings of the nodes in each of the route replies and selects the route with the highest metric.

A different approach is to provide incentive to nodes so that they comply with protocol rules, i.e., properly relay user data. The concept of fictitious currency is introduced in [6] to endogenize the behavior of the assumed greedy nodes, which would forward packets in exchange for currency. Each intermediate node purchases from its predecessor the received data packet and sells it to its successor along the path to the destination. Eventually, the destination pays for the received packet. This scheme assumes the existence of an overlaid geographic routing infrastructure and a PKI. All nodes are preloaded with an amount of currency, have unique identifiers, and are associated with a pair of private/public keys. Finally, the cryptographic operations related to the currency transfers are performed by a physically tamper-resistant module.

Another approach appropriate for MANET, which departs significantly from the two above-mentioned schemes, is presented below. Low-cost cryptography is used to protect the integrity and origin authenticity of exchanged data, without placing any overhead at intermediate nodes. Moreover, the feedback that determines the security of the chosen paths originates only from trusted destinations, thus allowing safe inferences on the quality of the paths. Finally, the reliability and fault tolerance of data transmissions is enhanced significantly.




Figure 21.3. SMT Protocol Uses Multiple Diverse Paths Connecting the Source and the Destination

In particular, the APS contains paths that have not been detected as failed, either due to path breakage or because of the presence of an adversary on the path.

21.6.1 Secure Message Transmission Protocol

The SMT protocol [18, 31] is a network-layer secure and fault-tolerant data-forwarding scheme, tailored to the MANET characteristics. In short, SMT is provided with routing information by a protocol such as SRP. This allows SMT to determine a set of diverse paths connecting the source and the destination, as shown in the example of Figure 21.3. Then, it introduces limited transmission redundancy across the paths, by dispersing a message into N pieces, so that successful reception of any M -out-of- N pieces allows the reconstruction of the original message at the destination. Each piece, transmitted over one path, is equipped with a cryptographic header that provides origin authentication, integrity, and replay protection. Upon reception of a number of pieces, the destination informs the source of which pieces, and thus routes, were intact. To enhance the robustness of the feedback mechanism, the small-sized acknowledgments, also protected by a cryptographic header, are maximally dispersed, so that successful reception of one piece is sufficient. If less than M pieces of the message were received, the source retransmits the remaining pieces over the intact routes or in general the ones deemed as more secure. If too few pieces were acknowledged or too many messages remain outstanding, the protocol adapts its operation by determining a different path set, reencoding undelivered messages, and reallocating pieces over the path set. Otherwise, it proceeds with subsequent message transmissions.

SMT exploits MANET features such as the topological redundancies, interoperates widely with accepted techniques such as on-demand route discovery and source routing, relies on a security association only between the source and the destination, and makes use of highly efficient symmetric-key cryptography. Moreover, the routing decisions are made by the querying node, based on the feedback that the destination and the underlying SRP provide. At the same time, no additional processing overhead is

imposed on intermediate nodes, which do not perform any cryptographic operation but simply relay the message pieces. However, the use of multiple paths and the resultant greater number of nodes involved in the forwarding of a single message can be admittedly considered as the price to pay to achieve the sought robustness.

On the one hand, SMT's robustness can be enhanced by the adaptation of parameters such as the number of paths and the ratio of the numbers of transmitted to required pieces, termed as the redundancy or dispersion factor. On the other hand, in a low-risk environment with limited malicious failures, the same parameters can be adjusted so that the imposed transmission overhead is reduced to a level close to that of a single-path scheme. An additional element that contributes to the flexibility of SMT is that different algorithms can be implemented for the selection of the path set, based on different metrics and interpretations of the network feedback. SMT can yield 100 percent successful message reception even in a highly adverse environment, when, for example, 20 percent of the network nodes are malicious, while keeping the message and computation overhead low.

The two communicating end nodes make use of the Active Path Set (APS), comprising diverse paths that are not deemed failed. The sender invokes the underlying route discovery protocol, updates its network topology view, and then determines the APS for a specific destination. This model can be extended to multiple destinations, with one APS per destination. At the receiver's side, the APS is used for the transmission of the feedback, but if links are not bidirectional, the destination will have to determine its own reverse APS.

The dispersion of messages, which is performed by the information dispersal algorithm (IDA) [22], is coupled to the APS characteristics through an appropriate selection of the dispersion algorithm parameters. For example, in low connectivity conditions (small number of disjoint paths), the sender may increase the redundancy factor to provide increased assurance and possibly low transmission delay. The adaptation of the protocol is the result of the interplay among the following parameters:

- K , the (sought) cardinality of APS
- k , the S, T -connectivity, i.e., the maximum number of S, T node-disjoint paths from the source (S) to the destination (T)
- r , the redundancy factor of the IDA encoding
- x , the maximum number of malicious nodes

Clearly, the condition for successful reception is $x \leq [K \cdot (1 - r^{-1})]$ which demonstrates the coupling among choices of parameters.

In particular, K can be determined as a function of r , so that the probability of successful transmission is maximized. (Note that K is equal to N




Figure 21.4. APS with Three Paths

For an APS with three paths, the source can disperse each message into three pieces and transmit them across APS. The destination responds to each message M_k with an acknowledgment ACK_k , notifying explicitly which pieces were received. This feedback allows the source to quickly update the rating of the APS paths and retransmit lost pieces across the operational paths if the message cannot be reconstructed at the destination.

when one message piece is allocated per path.) To do so, the source starts by determining an APS of the k shortest, in number of hops, node disjoint paths. Then, let P_{GOAL} be the sought probability of successful reconstruction of a dispersed message. P_{GOAL} can be provided from the application layer and may correspond to the features of the supported application, for example. Given P_{GOAL} and k , the node calculates the required redundancy factor, r_{GOAL} , for a given or estimated fraction of present adversaries. The source disperses outgoing messages with the redundancy value closest to r_{GOAL} , with M and N selected to minimize the transmission overhead.

Once dispersed, the message pieces are transmitted across APS. If $N < k$, the node selects the N paths of the APS with the highest rating. If the receiver cannot reconstruct the message, the source retransmits the pieces that were not received, according to the feedback provided by the destination. Message pieces are retransmitted by SMT a maximum number of times, $Retry_{MAX}$, which is a protocol selectable parameter. If all retransmissions fail, the message is discarded. This way, limited retransmissions enhance the efficiency of SMT by alleviating the overhead from retransmitting the entire amount of data while SMT promptly detects failures and adapts its operation to remain effective and efficient (Figure 21.4).

The transmission of data is continuous over the APS, with retransmissions placed at the head of the queue, upon reception of the feedback. The continuous usage of the APS allows SMT to update quickly its assessment on the quality of the paths. For each successful or failed piece, the rating of the corresponding path is increased or decreased, respectively. When the rating drops below a threshold, the path is discarded. The path rating is also decreased slowly as time goes by to reduce the chance of using a stale path. Moreover, the simultaneous routing over a number of paths, if not the entire APS, provides the opportunity for low-cost probing of the paths. In particular, the source can easily tolerate the loss of a piece that was transmitted over a low-rated path.

21.7 Discussion

The fast development of mobile ad hoc networking technology over the last few years, with satisfactory solutions to a number of technical problems, supports the vision of widely deployed mobile ad hoc networks with self-organizing features and without the necessity of a preexisting infrastructure. In this context, the secure operation of such infrastructureless networks becomes a primary concern. Nevertheless, the provision of security services is dependent on the characteristics of the supported application and the networked environment, which may vary significantly. At one extreme, we can think of a library or an Internet café, which provides short-range wireless connectivity to patrons, without any access constraint other than the location of the mobile device. At the other extreme, a military or public safety unit can make use of powerful mobile devices capable of performing expensive cryptographic operations. Such devices would communicate only with other trusted devices.

Between these two ends of the spectrum, a multitude of MANET instances will provide different services, assume different modes of interaction and trust models, and admit solutions such as the ones surveyed above. Moreover, it is probable that instead of a clear-cut distinction among network instances, devices and users with various security requirements will coexist in a large, open, frequently changing ubiquitous network.

In this context, an important related issue is the IP addressing scheme employed in the MANET environment. The common assumption that node credentials (e.g., certificates) are bound to IP addresses may need to be revisited, because one can imagine that roaming nodes will join MANET subdomains and IP addresses will be assigned dynamically (e.g., DHCP [8] or IPv6 auto-configuration [26]) or even randomly (e.g., Zero-Configuration [12]).

A type of ad hoc network with particular requirements is a sensor network, which requires multi-hop communication throughout a network of hundreds or even thousands of nodes, with relatively infrequent topological changes. It is expected that a single organization will undertake the deployment and administration of these networks. Moreover, sensing devices have limited computational capabilities, network transmission rates are relatively low, and communications are mostly data driven. These requirements may affect, in different ways, the design of security measures for sensor networks, as demonstrated by the schemes proposed in the literature.

One of the proposals to secure sensor networks provides a protocol for data authentication, integrity, and freshness and a lightweight implementa-

tion of an authenticated broadcast protocol [21]. The scheme targets a restricted, infrastructure-oriented environment, with a trusted central entity instituted by a set of base stations. Sensor nodes communicate only with a base station, which broadcasts messages toward the sensors. The base station and all nodes initially possess a symmetric encryption and authentication key, which secures the exchanged traffic, and later, the base station periodically broadcasts the key that was used to authenticate transmissions during the last period.

An approach that has similarities but targets a more general setting proposes a key management scheme for sensor networks [2]. The focus is on resource-constrained large sensor networks, comprising nodes that are assumed tamper-resistant and equipped with a secret group key. Similarly to the previous scheme, the use of symmetric key cryptography is proposed as the only feasible, low-cost solution. However, frequent rekeying (i.e., periodic regeneration of the single key that is used to encrypt all data transmitted by sensors) is proposed to protect it from possible compromise. To make this reconfiguration operation efficient, the sensors are organized into clusters with a two-hop diameter and clusterheads are elected and form a backbone. Then, from a subset of the backbone, a randomly elected node generates the new key.

The simplified trust models of sensor networks, which, nonetheless, lead to efficient solutions, may not necessarily be usable in other ad hoc networking instances. The circumstantial coexistence of disparate nodes or the requirement of fine-grained trust relationships call for solutions that can adapt to specific contexts and support the corresponding application. However, although the requirements of the application are expected to dictate the characteristics of the required security mechanisms, some aspects of security, such as confidentiality, may not be different at all in the MANET context. Instead, the greatest challenge is to safeguard the basic network operation.

In particular, the securing of the network topology discovery and data forwarding is a prerequisite for the secure operation of mobile ad hoc networks in any adverse environment. Furthermore, the protection of the functionality of the networking protocols will be in many cases orthogonal to the security requirements and the security services provided at the application layer. For example, a transaction can be secured when the two communicating end nodes execute a cryptographic protocol based on established mutual trust, with the adversary being practically unable to attack the protocol. But this does not imply that the nodes are secure against denial of service attacks; the adversary can still abuse the network protocols, and in fact, do it with little effort compared to the effort needed to compromise the cryptographic protocol.

The self-organizing networking infrastructure has to be protected against misbehaving nodes, with the use of low-cost cryptographic tools, under the least restrictive trust assumptions. Moreover, the overhead stemming from such security measures should be imposed mostly, if not entirely, on nodes that communicate in a secure manner and that directly benefit from these security measures. Furthermore, we believe that the salient MANET features and the unique operational requirements of these networks call for security mechanisms that are primarily present at, and closely interwoven with, the network-layer operation, to realize the full potential of this promising new technology.

References

1. D. Balfanz, D.K. Smetters, P. Stuart, and H.C. Wang, Talking to Strangers: Authentication in Ad Hoc Networks, *Proceedings of Network and Distributed System Security Symposium*, San Diego, CA, February 2002.
2. S. Basagni, K. Herrin, E. Rosti, and D. Bruschi, Secure Pebblenets, *Proceedings of 2nd MobiHoc*, Long Beach, CA, October 2001.
3. S.M. Bellovin and M. Merritt, Encrypted Key Exchange: Password-Based Protocols Secure Against Dictionary Attacks, *Proceedings of the IEEE Symposium on Security and Privacy*, Oakland, CA, May 1992.
4. Bluetooth Special Interest Group, Specifications of the Bluetooth System, <http://www.bluetooth.com>.
5. M. Brown, D. Cheung, D. Hankerson, J.L. Hernandez, M. Kirkup, and A. Menezes, PGP in Constrained Wireless Devices, *Proceedings of 9th USENIX Symposium*, Denver, August 2000.
6. L. Buttyan and J.P. Hubaux, Enforcing Service Availability in Mobile Ad Hoc WANs, *Proceedings of 1st MobiHoc*, Boston, August 2000.
7. B. Dahill, B.N. Levine, E. Royer, and C. Shields, A Secure Routing Protocol for Ad Hoc Networks, Technical Report UM-CS-2001-037, EE&CS, University of Michigan, Ann Arbor, August 2001.
8. R. Droms, Dynamic Host Configuration Protocol, IETF RFC 2131, March 1997.
9. L.M. Feeney, B. Ahlgren, and A. Westerlund, Spontaneous Networking: An Application-Oriented Approach to Ad Hoc Networking, *IEEE Communications Magazine*, June 2001, pp. 176–181.
10. M. Guerrero Zapata and N. Asokar, Securing Ad Hoc Routing Protocols, *Proceedings of the ACM Wireless Security Workshop (WiSe 2002)*, Atlanta, GA, 2002.
11. V. Gupta and S. Gupta, Securing the Wireless Internet, *IEEE Communications Magazine*, December 2001, pp. 68–74.
12. M. Hattig, Ed., Zero-Conf IP Host Requirements, Draft-ietf-zeroconf-reqts-09.txt, IETF MANET Working Group, August 2001.
13. J.P. Hubaux, L. Buttyan, and S. Capkun, The Quest for Security in Mobile Ad Hoc Networks, *Proceedings of 2nd MobiHoc*, Long Beach, CA, October 2001.
14. J. Kong, P. Zerfos, H. Luo, S. Lu, and L. Zhang, Providing Robust and Ubiquitous Security Support for Mobile Ad-Hoc Networks, *Proceedings of IEEE ICNP (International Conference on Network Protocols) 2001*, Riverside, CA, November 2001.
15. L. Lamport, Password Authentication with Insecure Communication, *Comm. of ACM*, vol. 24, pp. 770–772, 1981.
16. S. Marti, T.J. Giuli, K. Lai, and M. Baker, Mitigating Routing Misbehavior in Mobile Ad Hoc Networks, *Proceedings of 6th MobiCom*, Boston, August 2000.

17. P. Papadimitratos and Z.J. Haas, Secure Routing for Mobile Ad Hoc Networks, *Proceedings of SCS Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS 2002)*, San Antonio, TX, January 27–31, 2002.
18. P. Papadimitratos and Z.J. Haas, Secure Data Transmission in Mobile Ad Hoc Networks, *Proceedings of the ACM Wireless Security Workshop (WiSe 2003)*, San Diego, CA, 2003.
19. P. Papadimitratos and Z.J. Haas, Securing the Internet Routing Infrastructure, *IEEE Communications Magazine*, vol. 40, no. 10, October 2002.
20. C.E. Perkins, E.M. Royer, and S.R. Das, Ad hoc On-Demand Distance Vector Routing, *Draft-ietf-manet-aodv-08.txt*, IETF MANET Working Group, June, 2001.
21. A. Perrig, R. Szewczyk, V. Wen, D. Culler, and J.D. Tygar, SPINS: Security Protocols for Sensor Networks, *Proceedings of 7th Ann. Int'l. Conf. Mobile Computing and Networks (MobiCom 2001)*, Rome, 2001, pp. 189–199.
22. M.O. Rabin, Efficient Dispersal of Information for Security, Load Balancing, and Fault Tolerance, *Journal of ACM*, vol. 36, pp. 335–348, 1989.
23. R. Rivest, A. Shamir, and L. Adleman, A Method for Obtaining Digital Signatures and Public Key Cryptosystems, *Comm. of ACM*, vol. 21, pp. 120–126, 1978.
24. F. Stajano and R. Anderson, The Resurrecting Duckling: Security Issues for Ad Hoc Wireless Networks, *Proceedings of Security Protocols, 7th Int'l. Workshop, Lecture Notes in Computer Science*, vol. 1796, Berlin: Springer-Verlag, 2000.
25. F. Stajano, The Resurrecting Duckling — What Next? *Proceedings of Security Protocols, 8th Int'l. Workshop, Lecture Notes in Computer Science*, vol. 2133, Berlin: Springer-Verlag, 2001.
26. S. Thomson and T. Narten, IPv6 Stateless Address Autoconfiguration, IETF RFC 1971, www.ietf.org.
27. G.R. Wright and W. Stevens, *TCP/IP Illustrated, Volume 2: The Implementation*, Boston: Addison Wesley, 1997.
28. S. Yi, P. Naldurg, and R. Kravets, Security-Aware Ad-Hoc Routing for Wireless Networks, UIUCDCS-R-2001-2241 Technical Report, August 2001.
29. L. Zhou and Z.J. Haas, Securing Ad Hoc Networks, *IEEE Network Magazine*, November/December 1999.
30. IETF Mobil Ad-hoc Networks (MANET) Working Group, <http://www.ietf.org/html.charters/manet-charter.html>.
31. P. Papadimitratos and Z.J. Haas, Secure Message Transmission in Mobile Ad hoc Networks, *Elsevier Ad Hoc Networks Journal*, vol. 1, no. 1, July 2003.