


INTRUSION DETECTION IN WIRELESS AD HOC NETWORKS

AMITABH MISHRA, KETAN NADKARNI, AND ANIMESH PATCHA, VIRGINIA TECH



With so much advancement in hacking, if attackers try hard enough, they will eventually succeed in infiltrating the system. This makes it important to monitor what is taking place on a system and look for suspicious behavior. Intrusion Detection Systems do just that.

ABSTRACT

Intrusion detection has, over the last few years, assumed paramount importance within the broad realm of network security, more so in the case of wireless ad hoc networks. These are networks that do not have an underlying infrastructure; the network topology is constantly changing. The inherently vulnerable characteristics of wireless ad hoc networks make them susceptible to attacks, and it may be too late before any counter action can take effect. Second, with so much advancement in hacking, if attackers try hard enough they will eventually succeed in infiltrating the system. This makes it important to constantly (or at least periodically) monitor what is taking place on a system and look for suspicious behavior. Intrusion detection systems (IDSs) do just that: monitor audit data, look for intrusions to the system, and initiate a proper response (e.g., email the systems administrator, start an automatic retaliation). As such, there is a need to complement traditional security mechanisms with efficient intrusion detection and response. In this article we present a survey on the work that has been done in the area of intrusion detection in mobile ad hoc networks.

INTRODUCTION

Wireless ad hoc networks have been in focus within the wireless research community. Essentially, these are networks that do not have an underlying fixed infrastructure. Mobile hosts “join” on the fly and create a network on their own. With the network topology changing dynamically and the lack of a centralized network management functionality, these networks tend to be vulnerable to a number of attacks.

Mobile nodes within one another’s radio range can communicate through wireless links and thus dynamically form a network. Wireless devices that are not in direct range communicate via intermediate devices; this is known as multihop communication. Thus, an ad hoc network is a collection of autonomous nodes that form a dynamic purpose-specific multihop radio network in a decentralized fashion. The quintessential nature of such networks is the conspicuous absence of a fixed support infrastructure such as

mobile switching centers, base stations, access points, and other centralized machinery seen in traditional wireless networks. The network topology is constantly changing as a result of nodes joining in and moving out. Packet forwarding, routing, and other network operations are carried out by the individual nodes themselves.

Wireless ad hoc networks find application in military operations so that planes, tanks, and moving personnel can communicate. Rescue missions and emergency situations also find use for such networks. Other examples include virtual classrooms and conferences wherein people can set up a network on the spot through their laptops, PDAs, and other mobile devices, assuming they share the same physical medium such as direct sequence spread spectrum (DSSS) or frequency hopped spread spectrum (FHSS).

The unreliability of wireless links between nodes, constantly changing topology due to the movement of nodes in and out of the network, and lack of incorporation of security features in statically configured wireless routing protocols not meant for ad hoc environments all lead to increased vulnerability and exposure to attacks. Security in wireless ad hoc networks is particularly difficult to achieve, notably because of the limited physical protection of each node, the sporadic nature of connectivity, the absence of a certification authority, and the lack of a centralized monitoring or management unit. Intrusion prevention is not guaranteed to work all the time; this clearly underscores the need for intrusion detection as a frontline security research area under the umbrella of ad hoc network security. If an intrusion is detected quickly enough, the intruder can be identified and ejected from the system before any damage is done or any data are compromised. Moreover, an effective intrusion detection system (IDS) can serve as a deterrent, acting to prevent intrusions. Intrusion detection enables the collection of information about intrusion techniques that can be used to strengthen the intrusion prevention facility. In this article we look at how ad hoc networks to a certain extent can be secured using traditional techniques. We then examine the different intrusion detection techniques proposed for these networks.

The rest of the article is organized as follows.

We present the characteristics of wireless ad hoc networks that make them so vulnerable to attacks. The fundamentals of intrusion detection are covered along with a classification of these systems. We then look at the requirements and characteristics of IDSs. The *pièce de résistance* of this article presents a state-of-the-art view of the research in intrusion detection in the ad hoc environment. We devote a section to comparing different intrusion detection schemes against a set of attributes that are desirable in any such scheme.

SECURITY VULNERABILITIES IN MOBILE AD HOC NETWORKS

There are various reasons why wireless ad hoc networks are at risk, from a security point of view. We next discuss the characteristics, which make these networks vulnerable to attacks.

In traditional wireless networks, mobile devices associate themselves with an access point, which is in turn connected to other wireline machinery such as a gateway or name server that manages the network management functions. Ad hoc networks, on the other hand, do not have a centralized piece of machinery such as a name server, which if present as a single node can be a single point of failure. The absence of infrastructure and the subsequent absence of authorization facilities impede the usual practice of establishing a line of defense, distinguishing nodes as trusted and nontrusted. There may be no ground for an a priori classification, since all nodes are required to cooperate in supporting the network operation, while no prior security association (SA) can be assumed for all the network nodes. Freely roaming nodes form transient associations with their neighbors, joining and leaving subdomains independently with and without notice.

An additional problem related to the compromised nodes is the potential *Byzantine* failures encountered within mobile ad hoc network (MANET) routing protocols wherein a set of nodes could be compromised in such a way that incorrect and malicious behavior cannot be directly noted at all. Such malicious nodes can also create new routing messages and advertise nonexistent links, provide incorrect link state information, and flood other nodes with routing traffic, thus inflicting Byzantine failures on the system.

The wireless links between nodes are highly susceptible to link attacks, which include passive eavesdropping, active interfering, leakage of secret information, data tampering, impersonation, message replay, message distortion, and denial of service (DoS). Eavesdropping might give an adversary access to secret information, violating confidentiality. Active attacks might allow the adversary to delete messages, inject erroneous messages, modify messages, and impersonate a node, thus violating availability, integrity, authentication, and nonrepudiation.

The presence of even a small number of adversarial nodes could result in repeatedly compromised routes; as a result, the network nodes would have to rely on cycles of timeout and new

route discoveries to communicate. This would incur arbitrary delays before the establishment of a non-corrupted path, while successive broadcasts of route requests would impose excessive transmission overhead. In particular, intentionally falsified routing messages would result in DoS experienced by the end nodes.

Moreover, the battery-powered operation of ad hoc networks gives attackers ample opportunity to launch a DoS attack by creating additional transmissions or expensive computations to be carried out by a node in an attempt to exhaust its batteries.

Attacks against MANETs can be divided into two groups: Passive attacks typically involve only eavesdropping of data whereas active attacks involve actions performed by adversaries, for instance the replication, modification and deletion of exchanged data. External attacks are typically active attacks that are targeted to prevent services from working properly or shut them down completely. Intrusion prevention measures like encryption and authentication can only prevent external nodes from disrupting traffic, but can do little when compromised nodes internal to the network begin to disrupt traffic. Internal attacks are typically more severe attacks, since malicious insider nodes already belong to the network as an authorized party and are thus protected with the security mechanisms the network and its services offer. Thus, such compromised nodes, which may even operate in a group, may use the standard security means to actually protect their attacks.

In summary, a malicious node can disrupt the routing mechanism employed by several routing protocols in the following ways.

Attack the route discovery process by:

- Changing the contents of a discovered route
- Modifying a route reply message, causing the packet to be dropped as an invalid packet
- Invalidating the route cache in other nodes by advertising incorrect paths
- Refusing to participate in the route discovery process

Attack the routing mechanism by:

- Modifying the contents of a data packet or the route via which that data packet is supposed to travel
- Behaving normally during the route discovery process but drop data packets causing a loss in throughput

Generate false route error messages whenever a packet is sent from a source to a destination.

Launch DoS attacks by:

- Sending a large number of route requests. Due to the mobility aspect of MANETs, other nodes cannot make out whether the large number of route requests are a consequence of a DoS attack or due to a large number of broken links because of high mobility.
- Spoofing its IP and sending route requests with a fake ID to the same destination, causing a DoS at that destination.

The above discussion makes it clear that ad hoc networks are inherently insecure, more so than their wireline counterparts, and need intrusion detection schemes before it is too late to counter an attack. If there are attacks on a sys-

The wireless links between nodes are highly susceptible to link attacks, which include passive eavesdropping, active interfering, leakage of secret information, data tampering, impersonation, message replay, message distortion, and denial of service.

An IDS is generally controlled by the configuration settings that would specify how and where to collect audit data, how to respond to intrusions, etc. Access to these configuration settings would give a potential intruder vital information on which avenues of attack are likely to go undetected.

tem, one would like to detect them as soon as possible (ideally in real time) and take appropriate action. This is essentially what an IDS does. We now discuss the basics of an IDS and provide a classification of such systems.

INTRUSION DETECTION SYSTEMS: A BRIEF OVERVIEW

Intrusion detection can be defined as the automated detection and subsequent generation of an alarm to alert the security apparatus at a location if intrusions have taken place or are taking place. An IDS is a defense system that detects hostile activities in a network and then tries to possibly prevent such activities that may compromise system security. IDSs achieve detection by continuously monitoring the network for unusual activity. The prevention part may involve issuing alerts as well as taking direct preventive measures such as blocking a suspected connection. In other words, intrusion detection is a process of identifying and responding to malicious activity targeted at computing and networking resources. In addition, IDS tools are capable of distinguishing between insider attacks originating from inside the network and external ones. Unlike firewalls which are the first line of defense, IDSs come into the picture only after an intrusion has occurred and a node or network has been compromised. That is why IDSs are aptly called the second line of defense.

Generally speaking, an IDS:

- Is NOT an antivirus program designed to detect malicious software's such as viruses, Trojans, and worms.
- Is NOT a network logging system used, for example, to detect complete vulnerability to any DoS attack across a congested network. These are network traffic monitoring systems.
- Is NOT a vulnerability assessment tool that checks for bugs and flaws in operating systems and network services. Such an activity would fall under the purview of security scanners.

A basic model of an IDS is likely to include quite a few elements. Primarily, intrusion detection decisions are based on collected audit data. Sources of data can include keyboard input, command-based logs, and application-based logs. Audit data is stored either indefinitely, for later reference, or temporarily, awaiting processing. The humongous volume of data makes this a crucial element in an IDS. One or many algorithms are executed to find evidence in the audit trail of suspicious behavior. An IDS is generally controlled by the configuration settings that would specify how and where to collect audit data, how to respond to intrusions, and so on. Access to these configuration settings would give a potential intruder vital information on which avenues of attack are likely to go undetected. Reference data stores information about known intrusion signatures (for misuse systems) or profiles of normal behavior (for anomaly systems). The processing element must frequently store intermediate results, an example of which might be information about partially fulfilled intrusion signatures. The space needed to store this active data can grow quite large too. And finally, the alarm part

of the system handles all output from the system. Examples include automated response to suspicious activity and notification of the user.

Intrusion detection can be classified into three broad categories: anomaly detection, signature or misuse detection, and specification-based detection. We discuss each of these as per the taxonomy proposed in [1].

Anomaly detection: In an anomaly detection system a baseline profile of normal system activity is created. Any system activity that deviates from the baseline is treated as a possible intrusion. The problems with strict anomaly detection are that:

- Anomalous activities that are not intrusive are flagged as intrusive.
- Intrusive activities that are not anomalous result in false negatives.

One disadvantage of anomaly detection for mobile computing is that the normal profile must be periodically updated and the deviations from the normal profile computed. The periodic calculations can impose a heavy load on some resource constrained mobile devices; perhaps a lightweight approach that involves comparatively less computation might be better suited.

Misuse detection: In misuse detection, decisions are made on the basis of knowledge of a model of the intrusive process and what traces it ought to leave in the observed system. *Legal* or *illegal* behavior can be defined and observed behavior compared accordingly. Such a system tries to detect evidence of intrusive activity irrespective of any knowledge regarding the background traffic (i.e., the normal behavior of the system).

Specification-based detection: Specification-based detection defines a set of constraints that describe the correct operation of a program or protocol, and monitors the execution of the program with respect to the defined constraints. This technique may provide the capability to detect previously unknown attacks, while exhibiting a low false positive rate.

An offshoot to misuse and anomaly detection is compound detection, which is basically a misuse inspired system that forms a compound decision in view of a model of both the normal behavior of the system and the intrusive behavior of the intruder. The detector operates by detecting the intrusion against the background of the normal traffic in the system. These detectors have a much better chance of correctly detecting truly interesting events in the supervised system, since they both know the patterns of intrusive behavior and can relate them to the normal behavior of the system. They would at the very least be able to qualify their decisions better.

INTRUSION RESPONSE

The type of intrusion response for wireless ad hoc networks depends on the type of intrusion, the network protocols and applications in use, and the confidence (or certainty) in the evidence. A few likely responses include:

- Reinitializing communication channels between nodes (e.g., force rekey).
- Identifying the compromised nodes and reorganizing the network to preclude the compromised nodes.

- The IDS agent informing the end user, who may in turn do his/her own investigation and take appropriate action.
- Initiating a re-authentication request to all nodes in the network to prompt the end users to authenticate themselves (and hence their wireless nodes) using out-of-band mechanisms (like visual contacts). Only the re-authenticated nodes, which may collectively negotiate a new communication channel, will recognize each other as legitimate. That is, the compromised/malicious nodes can be excluded.

REQUIREMENTS FOR AN INTRUSION DETECTION SYSTEM FOR MOBILE AD HOC NETWORKS

There are two key requirements that any IDS must fulfill. These are *effectiveness* — how to make the intrusion detection system classify malign and benign activity correctly — and *efficiency* — how to run the IDS in a cost effective manner as far as possible. In other words, these two requirements in essence suggest that an IDS should detect a substantial percentage of intrusions into the supervised system, while keeping the false alarm rate at an acceptable level at a lower cost. It is expected that an ideal IDS is likely to support several of the following requirements:

- The IDS should not introduce a new weakness in the MANET. That is, the IDS itself should not make a node any weaker than it already is.
- An IDS should run continuously and remain transparent to the system and users.
- The IDS should use as little system resources as possible to detect and prevent intrusions. IDSs that require excessive communication among nodes or run complex algorithms are not desirable.
- It must be fault-tolerant in the sense that it must be able to recover from system crashes, hopefully recover to the previous state, and resume the operations before the crash
- Apart from detecting and responding to intrusions, an IDS should also resist subversion. It should monitor itself and detect if it has been compromised by an attacker.
- An IDS should have a proper response. In other words, an IDS should not only detect but also respond to detected intrusions, preferably without human intervention.
- Accuracy of the IDS is another major factor in MANETs. Fewer false positives and false negatives are desired.
- It should interoperate with other intrusion detection systems to collaboratively detect intrusions. For example, the Internet Engineering Task Force (IETF) Intrusion Detection Working Group (IDWG) [2] is working toward proposing such a specification.

INTRUSION DETECTION IN MANETS

Quite a bit of research work has already been done in intrusion detection for traditional wired networks. However, applying the research of wired networks to wireless networks is not an easy plug-and-play task because of key architec-

tural differences, principal among them being the lack of fixed infrastructure. The absence of physical infrastructure facilitates the attacker's task since it is easier to eavesdrop on network traffic in a wireless environment.

Wireless ad hoc networks, due to their vulnerabilities, provide a tougher challenge for designing an IDS. Without centralized audit points such as routers and gateways, an IDS for ad hoc networks is limited to using only the current traffic coming in and out of the node as audit data. Another key requirement is that the algorithms the IDS uses must be distributed in nature, and should take into account the fact that a node can only see a portion of the network traffic. Moreover, since ad hoc networks are dynamic and nodes can move about freely, there is a possibility that one or more nodes could be captured and compromised, especially if the network is in a hostile environment. If the algorithms of the IDS are cooperative, it becomes important to be skeptical of which nodes one can trust. Therefore, intrusion detection systems on ad hoc networks have to be wary of attacks made from nodes in the network itself, not just attacks from outside the network. Also, mobile networks cannot communicate as frequently as their wired counterparts to detect intrusions in order to conserve bandwidth resources. Bandwidth and other issues such as battery life compound the problem even further. The availability of partial audit data makes it harder to distinguish an attack from regular network use.

In this section we present a state-of-the-art view of research in IDSs for MANETs, including proposed architectures and development work that is going on.

A DISTRIBUTED IDS

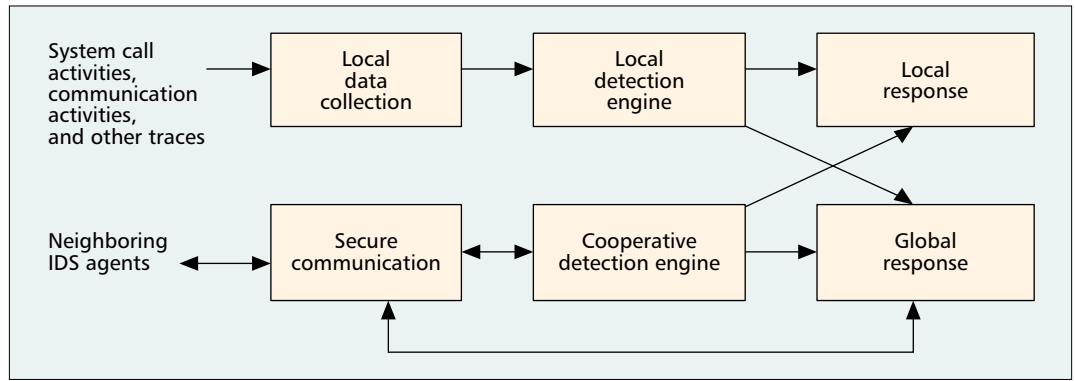
In their pioneering work on intrusion detection in MANETs, Zhang and Lee describe a distributed and cooperative intrusion detection model where every node in the network participates in intrusion detection and response [3]. In this model, an IDS agent runs at each mobile node, and performs local data collection and local detection, whereas cooperative detection and global intrusion response can be triggered when a node reports an anomaly. The authors consider two attack scenarios separately:

- Abnormal updates to routing tables
- Detecting abnormal activities in layers other than the routing layer

The internals of an IDS agent are structured into six pieces, as shown in Fig. 1. Each node does local intrusion detection independently, and neighboring nodes collaboratively work on a larger scale. Individual IDS agents placed on each and every node run independently and monitor local activities (including user, systems, and communication activities within the radio range), detect intrusions from local traces, and initiate responses. Neighboring IDS agents cooperatively participate in global intrusion detection actions when an anomaly is detected in local data or if there is inconclusive evidence. The data collection module gathers local audit traces and activity logs that are used by

Quite a bit of research work has already been done in intrusion detection for traditional wired networks. However, applying the research of wired networks to wireless networks is not an easy plug-and-play task because of key architectural differences, principal among them being the lack of fixed infrastructure.

The data collection module gathers local audit traces and activity logs that are used by the local detection engine to detect local anomaly. Detection methods that need broader data sets or that require collaborations among local IDS agents use the cooperative detection engine.



■ **Figure 1.** An intrusion detection system for MANETs.

the local detection engine to detect local anomaly. Detection methods that need broader data sets or require collaborations among local IDS agents use the cooperative detection engine. Both the local and global response modules provide intrusion response actions. The local response module triggers actions local to this mobile node (e.g., an IDS agent alerting the local user), while the global one coordinates actions among neighboring nodes, such as the IDS agents in the network electing a remedial action. A secure communication module provides a high-confidence communication channel among IDS agents.

The main contribution of [3] is that it presents a distributed and cooperative intrusion detection architecture based on statistical anomaly detection techniques. This article was among the first that had such a detailed distributed design. The design of actual detection techniques, their performance as well as verification, however, were not addressed in the article.

AODV PROTOCOL-BASED IDS

Bhargava *et al.* [4] proposed an intrusion detection and response model (IDRM) to enhance security in the Ad Hoc On Demand Distance Vector (AODV) routing protocol [5]. The intrusion detection model proposed by the authors is an extension of the model described above.

Figure 2 illustrates how the IDRM provides security to AODV. In this scheme, each node employs the IDRM that utilizes neighborhood information to detect misbehavior of its neighbors. When the misbehavior count for a node exceeds a predefined threshold, the information is sent out to other nodes as part of global response. The other nodes receive this information, check their local *Malcount* for this malicious node, and add their results to the initiator's response. In the intrusion response model (IRM), a node identifies that another node has been compromised when its *Malcount* increases beyond the threshold value for that allegedly compromised node. In such cases, it propagates this information to the entire network by transmitting a special type of packet called a *MAL* packet. If another node also suspects that the detected node is compromised, it reports its suspicion to the network and retransmits another special type of packet called *REMAL*. If two or more nodes report about a particular node,

another special packet, called a *PURGE* packet, is transmitted to isolate the malicious node from the network. All nodes that have a route through the compromised node look for newer routes. All packets received from a compromised node are dropped.

Some of the internal attacks include distributed false route request, DoS, impersonation, and compromise of a destination. The authors have proposed to identify these internal attacks in the following ways:

Distributed false route request: A malicious node might send frequent unnecessary route requests. When the nodes in the network receive a number of route requests greater than a threshold count by a specific source for a destination in a particular time interval, the node is declared malicious.


Denial of service: A malicious node launches the DoS attack by transmitting false control packets and using all the network resources. DoS can be launched by transmitting false routing messages or data packets. It can be identified if a node is generating control packets that are more than the threshold count in a particular time interval.

Destination is compromised: This attack is identified when the source does not receive a reply from the destination in a particular time interval. The neighbors generate probe/hello packets to determine connectivity.

Impersonation: It can be avoided if the sender encrypts the packet with its private key and other nodes decrypt with the public key of the sender. If the receiver is not able to decrypt the packet, the sender might not be the real source; hence, the packet is dropped.

TECHNIQUES FOR INTRUSION-RESISTANT AD HOC ROUTING ALGORITHMS

Techniques for Intrusion-Resistant Ad Hoc Routing Algorithms (TIARA) is a set of design techniques that strengthen MANETs against DoS attacks [6]. The TIARA mechanisms limit the damage sustained by MANETs from intrusion attacks and allow continued network operation at an acceptable level during such attacks. It provides protection against attacks on control routing traffic as well as data traffic, thereby providing a comprehensive defense against intruders. Because of routing algorithm indepen-



■ **Figure 2.** Handling of internal attacks.

dence it allows widespread applicability and supports secure enclaves for dynamic coalitions.

Research efforts at Architecture Technology Corporation are aimed at demonstrating a set of innovative design techniques, collectively called TIARA, that secure ad hoc networks against DoS attacks. The TIARA approach involves fully distributed lightweight firewalls for ad hoc wireless networks, distributed traffic policing mechanisms, intrusion-tolerant routing, distributed intrusion detection mechanisms, flow monitoring, reconfiguration mechanisms, multipath routing, and source-initiated route switching. The flow-based route access control (FRAC) rules define admissible flows. Per-flow security association is instantiated by secure session setup signaling protocol and contains information for packet authentication. Also, fast authentication enables low-overhead integrity checks on packet flow-ids and sequence numbers. There is referral-based resource allocation, which limits networks' exposure to resource usurpation by spurious sessions, and flows are assigned an initial allowable resource usage. Moreover, additional resources are only granted if the source of the flow can present referrals from a certain number of trusted nodes. Referrals have time-bound validity. Flow-specific sequence numbers limit and contain the impact of traffic replay attacks; sequence numbers are embedded within secret locations within each packet. The destination of flow monitors select flow parameters to detect intrusion-induced path failures, and multipath routing and source-initiated route switching divert flow through available alternate paths to circumvent intruders. Efforts are on to implement dynamic on-the-fly modifications to FRAC (firewall) policies, real-time referral-based resource allocation, lightweight implementation of traffic policing, fast authentication mechanisms resistant to traffic analysis, and embedding sequence numbers and path labels in encrypted packets. Although the proposed architecture seems to cover most of the important aspects of intrusion detection and prevention in MANETs, implementation of such a design methodology entails extensive modification of the routing algorithms in a MANET. A summary of countermeasures used in TIARA against intrusion attacks is shown in Table 1.

WATCHDOG-PATHRATER APPROACH

Sergio Marti *et al.* discussed two techniques that improve throughput in MANETs in the presence of compromised nodes that agree to forward packets but fail to do so [7]. A node may misbehave because it is overloaded, selfish, malicious, or broken. An overloaded node lacks the CPU

Intrusion attacks → Countermeasures ↓	Spurious traffic	Packet replay	Session flooding	Flow disruption	Route hijacking
FRAC	X				
Fast authentication	X	X	X		
Sequence numbers		X			
Referrals			X		
Flow monitoring				X	X
Multipath routing				X	X
Source init route switching				X	X


■ **Table 1.** A summary of TIARA countermeasures against intrusion attacks.

cycles, buffer space, or available network bandwidth to forward packets. A selfish node is unwilling to spend battery life, CPU cycles, or available network bandwidth to forward packets not of direct interest to it, even though it expects others to forward packets on its behalf. A malicious node launches a DoS attack by dropping packets. A broken node might have a software fault that prevents it from forwarding packets.


To mitigate the decrease in the throughput due to the above node categories, the authors use watchdogs that identify misbehaving nodes and a pathrater that helps routing protocols avoid these nodes. When a node forwards a packet, the node's watchdog verifies that the next node in the path also forwards the packet. The watchdog does this by listening promiscuously to the next node's transmissions. If the next node does not forward the packet, it is misbehaving. The watchdog detects misbehaving nodes. Every time a node fails to forward the packet, the watchdog increments the failure tally. If the tally exceeds a certain threshold, it determines that the node is misbehaving; this node is then avoided using the pathrater. The pathrater, run by each node in the network, combines knowledge of misbehaving nodes with link reliability data to pick the route most likely to be reliable. Each node maintains a rating for every other node it knows about in the network. It calculates a path metric by averaging the node ratings in the path.

The watchdog technique has its own advantages and weaknesses. Dynamic source routing (DSR) [8] with the watchdog has the advantage that it can detect misbehavior at the forwarding level, not just at the link level. Watchdog's weaknesses are that it might not detect a misbehaving node in the presence of:

Every node is responsible for detecting signs of intrusion locally and independently by monitoring activities such as user and system activities and the communication activities within the radio range, but neighboring nodes can collaboratively investigate in a broader range.



■ **Figure 3.** Node A does not hear B forward packet 1 to C, because B's transmission collides at A with packet 2 from source S.



■ **Figure 4.** Node A believes that B has forwarded packet 1 to C, although C never received the packet due to a collision with packet 2.

- Ambiguous collisions: These prevent node A from overhearing the transmission from node B, as shown in Fig. 3.
- Receiver collisions: Node A can only tell whether B has sent a packet, but not if node C received it or not, as shown in Fig. 4.
- Limited transmission power: A misbehaving node could limit its transmission power such that the signal is strong enough to be overheard by the previous node but too weak to be received by the true recipient.
- False misbehavior: This occurs when a node falsely reports other nodes as misbehaving.
- Partial dropping: A node can circumvent the watchdog by dropping packets at a lower rate than the watchdog's configured minimum misbehaving threshold.

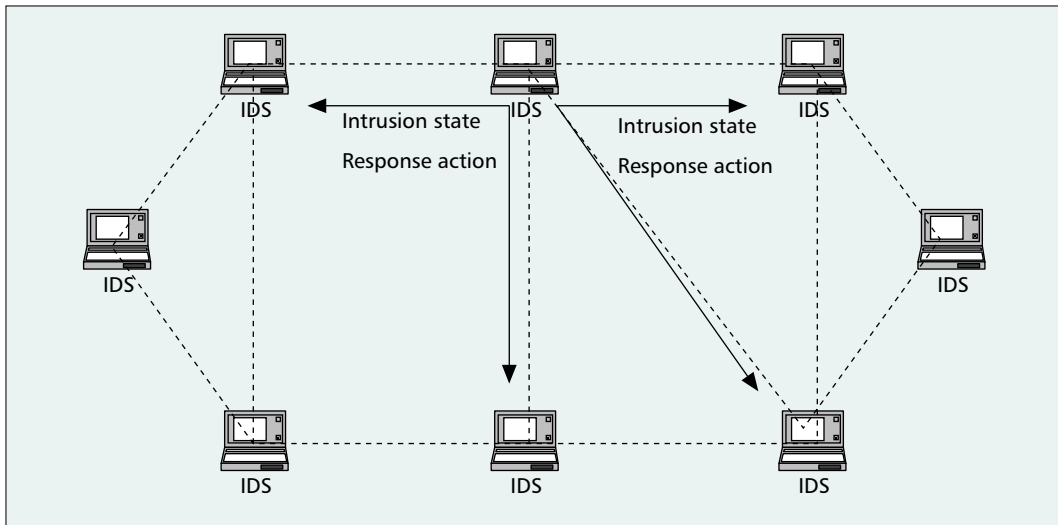
ANOMALY DETECTION FOR MOBILE WIRELESS NETWORKS

An anomaly detection architecture that was proposed in [9] is shown in Fig. 5. In this scheme every node in the MANET participates in intrusion detection and response. Every node is responsible for detecting signs of intrusion locally and independently by monitoring activities such as user and system activities and the communication activities within the radio range, but neighboring nodes can collaboratively investigate over a broader range. The internals of the detection scheme are conceptually shown in Fig. 1. Information-theoretic measures [10] such as entropy and conditional entropy are used to describe the characteristics of normal information flows, and classification algorithms are used to build anomaly detection models. For example, a classifier trained using normal data can be used to predict what normally the next event is given the previous n events. In monitoring, when the actual event is not what the classifier has predicted, there is an anomaly. When constructing a classifier, features with high information gain (or reduction in entropy) are needed. That is, a classifier needs feature value tests to partition the original (mixed and high entropy) dataset into pure (and low entropy) subsets, each ideally with one (correct) class of data.

Using the above mentioned framework, the following procedure is utilized for anomaly detection:

- Select (or partition) audit data so that the normal dataset has low (conditional) entropy.
- Perform appropriate data transformation according to the entropy measures (e.g., constructing new features with high information gain).
- Compute the classifier using training data.
- Apply the classifier to test data.
- Post-process alarms to produce intrusion reports.

Local routing information, including cache entries and traffic statistics, is used as an audit data source because remote nodes can be compromised and their data cannot be trusted. Since classifiers are used as detectors, there is a need to select and/or construct features from the available audit data that have high information gain. An unsupervised method is used to construct the feature set. First, a large feature set is constructed to cover a wide range of behaviors. Then a small number of training runs can be performed with the whole set of features on small audit data traces randomly chosen from previously stored audit logs. For each training run, a corresponding model is built. The features that appear in the models and have weights not smaller than a minimum threshold are selected into the essential feature set. For different routing protocols and different scenarios, the essential feature set is different. In practice, the feature set needs to be updated after a certain period, as the characteristics of routing behavior can change with time. The heuristic is that with sufficiently high dimension, data can be separated by a hyperplane, thus achieving the goal of classification. Given an execution trace, a detector is first applied to examine each observation. Then a post-processing scheme is used to examine the predictions and generate intrusion reports. A detection model can make spurious errors, and these false alarms should be filtered out. In contrast, a true intrusion session has "locality" (i.e., it tends to result in many alarms within a short time window). Therefore, these alarms can be grouped into a single intrusion report.



■ **Figure 5.** *IDS architecture for a wireless ad hoc network.*

The collaboration among the nodes is achieved using two types of data: security data to obtain complementary information from collaborating hosts, and intrusion alerts to inform others of a locally detected intrusion.

MOBILE AGENTS FOR INTRUSION DETECTION AND RESPONSE IN MANETS

Mobile agents are special kinds of agent that have the ability to move through large networks. In moving, the agents can interact with nodes, collect information, and execute tasks assigned to them. Mobile agents offer several advantages such as reduction in network load as well as latency, which is achieved by eliminating the need to move large amounts of data through the network via moving the analysis programs closer to the audit data. When portions of an IDS get destroyed or separated due to network partitioning, the mobile agents can still continue to work, thereby increasing the fault tolerance level of the network. The mobile agents tend to be independent of platform architectures, and thus enable agent-based IDSs to run under different operating system environments.

Local Intrusion Detection System (LIDS) — The LIDS is distributed in nature and utilizes mobile agents on each of the nodes of the ad hoc network [11]. In order to make local intrusions a global concern for the entire network, the LIDSs existing on different nodes collaborate. Collaboration among the nodes is achieved using two types of data: security data to obtain complementary information from collaborating hosts, and intrusion alerts to inform others of a locally detected intrusion. LIDS has chosen to use Simple Network Management Protocol (SNMP) data located in management information bases (MIBs) as the audit source because SNMP offers several advantages; principal is that the cost of local information collection is negligible if an SNMP agent is running on a node. Mobile agents (which must be autonomous and adaptive) are used to transport SNMP requests to remote hosts to overcome the unreliability of SNMP message transfer over UDP. A LIDS can delegate a specific mission to an agent that it will achieve in an autonomous and asynchronous manner without any help from its LIDS.

The LIDS architecture is shown in Fig. 6. The key elements of the architecture are:

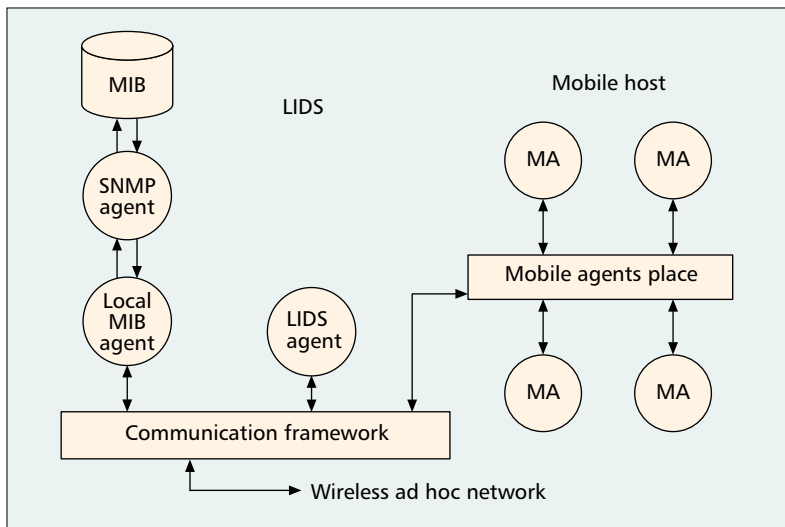
◇ A common communication framework to facilitate all external and internal communication with a LIDS

◇ Several data collecting agents for different tasks, such as:

- A local LIDS agent is in charge of local intrusion detection and response, and also for reacting to intrusion alerts provided by other nodes in order to protect itself against this intrusion.
- Mobile agents that collect and process data on remote hosts with an ability to transfer the results of a computation back to their home LIDS or to migrate to another node for further investigation. The mobile agent place is responsible for security control of these agents, but an agent should also be able to protect itself from malicious mobile agent places.
- The local MIB agent provides a means of collecting MIB variables for either mobile agents or the local LIDS agent. If SNMP runs on the node, the local MIB agent will be the interface with the running SNMP agent. For other scenarios an SNMP-based agent has to be developed to allow optimized updates and retrieval of the MIB variables used by intrusion detection. The local MIB agent would in that case act as an interface between the LIDS and this tailor-made agent.

In this design the local LIDS agent could use either misuse or anomaly detection as an intrusion detection mechanism. As far as response is concerned, as soon as a LIDS detects an intrusion locally it informs the other nodes of the network. Locally, the node is empowered to refuse connections with the suspicious node, exclude it when performing cooperative actions, or exclude it from its community until it re-authenticates itself. By being informed of intrusions on remote hosts, a LIDS can act as a security tool and prevent the intruder from attacking it. The authors recommend that for the best security in an ad hoc network, all the LIDSs on nodes should run and cooperate continuously.

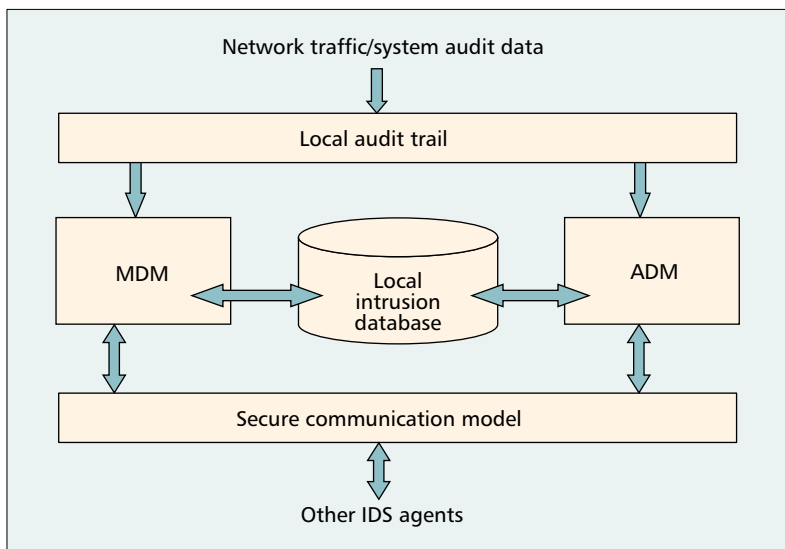
The novelty of this scheme lies in its use of SNMP data located in MIBs as audit sources



■ Figure 6. *LIDS architecture.*

and the use of mobile agents to process data at the source to reduce communication overheads. The use of a standard alert format, Intrusion Detection Message Exchange Format (IDMEF), and a protocol for transporting such alerts, Intrusion Detection Exchange Protocol (IDXP), ensures that IDSs running on a broad range of platforms can still interact and exchange intrusion-related information. On the downside, the authors do not consider compromised nodes broadcasting false intrusion-related information to the network.

Intrusion Detection Architecture Based on a Static Stationary Database — A distributed IDS has been proposed at Mississippi State University in which each node on the network has an IDS agent running on it [12]. The IDS agents on each node in the network work together via a cooperative intrusion detection algorithm to decide when and how the network is being attacked. The architecture is divided into parts: the mobile IDS agent, which resides on each node in the network, and the stationary secure database, which



■ Figure 7. *A proposed IDS based on a stationary secure database.*

contains global signatures of known misuse attacks and stores patterns of each user's normal activity in a non-hostile environment.

Mobile IDS Agents — Each node in the network will have an IDS agent running on it all the time. This agent is responsible for detecting intrusions based on local audit data and participating in cooperative algorithms with other IDS agents to decide if the network is being attacked. Each agent has five parts: a local audit trail, a local intrusion database (LID), a secure communication module, anomaly detection modules (ADMs), and misuse detection modules (MDMs), as shown in Fig. 7.

The LID is a local database that warehouses all information necessary for the IDS agent, such as the signature files of known attacks, the established patterns of users on the network, and the normal traffic flow of the network. The ADMs and MDMs communicate directly with the LID to determine if an intrusion is taking place.

The secure communication module is necessary to enable an IDS agent to communicate with other IDS agents on other nodes. It will allow the MDMs and ADMs to use cooperative algorithms to detect intrusions. It may also be used to initiate a global response when an IDS agent or a group of IDS agents detects an intrusion. Data communicated via the secure communication module needs to be encrypted.

The ADMs are responsible for detecting a different type of anomaly. There can be from one to many ADMs on each mobile IDS agent, each working separately or cooperatively with other ADMs.

The MDMs identify known patterns of attacks that are specified in the LID. Like the ADMs, if the audit data available locally is sufficient to determine if an intrusion is taking place, the proper response can be initiated. It is also possible for an MDM to use a cooperative algorithm to identify an intrusion.

Stationary Secure Database — The stationary secure database (SSD) acts as a secure trusted repository for mobile nodes to obtain information about the latest misuse signatures and find the latest patterns of normal user activity. It is assumed that the attacker will not compromise the SSD, as it is stored in an area of high physical security. The mobile IDS agents will collect and store audit data (user commands, network traffic, etc.) while in the field, and will transfer this information when they are attached to the SSD. The SSD will then use this information for data mining of new anomaly association rules. The SSD will also be the place where the system administrator can specify the newest misuse signatures. When the IDS agents are connected to SSD, they will gain access to the latest attack signatures automatically. Using the SSD to communicate the new attack signatures and establish new patterns of normalcy limits the amount of communication that must take place between IDS agents in the MANET. Despite all the benefits of having an SSD in a mobile IDS architecture, there are a few disadvantages in relying on a stationary database to provide vital IDS information. If an SSD is used, mobile nodes will have

to be attached to the non-mobile database periodically to stay up to date with the latest intrusion information. This may not be an option for some mobile ad hoc environments. Also, since the SSD must be a trusted source, it cannot be taken onsite without significant risk.

Distributed Intrusion Detection Using Mobile Agents — Kachirski and Guha have proposed a distributed intrusion detection system for ad hoc wireless networks based on mobile agent technology [13]. By efficiently merging audit data from multiple network sensors, their bandwidth-conscious scheme analyzes the entire ad hoc wireless network for intrusions at multiple levels, tries to inhibit intrusion attempts, and provides a lightweight low-overhead mechanism based on the mobile agent concept.

There is an efficient distribution of mobile agents with specific IDS tasks according to their functionality across a wireless ad hoc network. The agents used are dynamically updateable, have limited functionality, and can be viewed as components of a flexible, dynamically configurable IDS. Additionally, this scheme restricts computation-intensive analysis of overall network security state to a few key nodes. These nodes are dynamically elected, and overall network security is not entirely dependent on any particular node. The modular approach taken has advantages such as increased fault tolerance, communications cost reduction, improved performance of the entire network, and scalability.

The proposed IDS is built on a mobile agent framework as shown in Fig. 8. It is a nonmonolithic system and employs several sensor types that perform specific functions, such as:

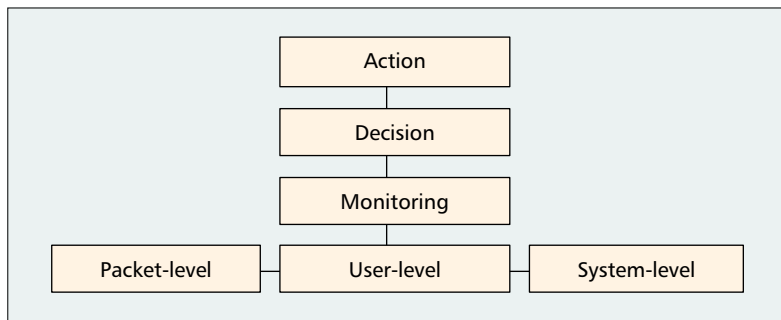
Network monitoring: Only certain nodes have sensor agents for network packet monitoring to preserve total computational power and battery power of mobile hosts.

Host monitoring: Every node on the MANET is monitored internally by a host monitoring agent. This includes monitoring system-level and application-level activities.

Decision-making: Every node decides on its intrusion threat level on a host-level basis. Certain nodes collect intrusion information and make collective decisions about the network-level intrusions.

Action: Every node has an action module responsible for resolving intrusion situations on a host.

There are three major agent categories: monitoring, decision-making, and action agents. Some are present on all mobile hosts, while others are distributed to only a selected group of nodes. While all nodes accommodate host-based monitoring sensors of the IDS, a distributed algorithm is utilized to assign a few nodes to host sensors that monitor network packets and agents that make decisions. The mobile network is logically divided into clusters with a single cluster head for each cluster that monitors packets within the cluster. The selected nodes host network monitoring sensors that collect all packets within the communication range and analyze the packets for known patterns of attacks. Monitoring agents are categorized into packet monitoring sensors, user activity sensors, and



■ Figure 8. Modular intrusion detection architecture.

system-level sensors. Local detection agents are located on each node of an ad hoc network, and act as user-level and system-level anomaly-based monitoring sensors. These agents look for suspicious activities on the host node, such as unusual process memory allocations, CPU activity, I/O activity, user operations (invalid login attempts with a certain pattern, super-user actions, etc.). If an anomaly is detected with strong evidence, a local detection agent will terminate the suspicious process or lock out a user and initiate reissue of security keys for the entire network. If some inconclusive anomalous activity is detected on a host node by a monitoring agent, the node is reported to the decision agent of the same cluster of which the suspicious node is a member. If more conclusive evidence is gathered about this node from any source (including packet monitoring results from a network monitoring agent), the action is undertaken by the agent on that node.

Decision agents are located on the same nodes as packet monitoring agents. A decision agent contains a state machine for all the nodes within the cluster it resides in. As intrusion or anomalous activity evidence is gathered for each node, the agent can decide that a node has been compromised by looking at reports from the node's own local monitoring agents and the packet monitoring information pertaining to that node. When a certain level of threat is reached for a node in question, the decision agent dispatches a command that an action must be undertaken by the local agents on that node. In time, the threat level decreases for each node in the decision agent's database.

SUMMARY

In this article we survey several intrusion detection schemes that have been proposed recently. The highlighted features of these schemes are summarized in Table 2. Severe memory constraints on a mobile device imply that misuse detection systems that need to store attack signatures will be relatively difficult to build and are likely to be less effective. Therefore, distributed anomaly detection is by far the methodology of choice for intrusion detection in MANETs; Table 2 clearly makes that point.

In Table 3 we compare different IDSs presented in this article against the attributes of an ideal IDS. These attributes are fault tolerance, scalability, interoperability with other IDSs, ability to detect new attack patterns, and whether

Proposed system	Highlighting features	Methodology
Distributed intrusion detection system for ad hoc networks	Statistical anomaly detection to detect local and global intrusions.	Distributed anomaly detection
Intrusion detection and response model for the AODV protocol	A collaborative threshold-based scheme where nodes watch neighbors for malicious activity. If two or more nodes report about a particular node, the malicious node is isolated from the network.	Distributed anomaly detection
Techniques for Intrusion-Resistant Ad Hoc Routing Algorithms (TIARA)	Presents routing-algorithm-independent general design principles and techniques that can be incorporated in MANETs for robust fault-tolerant networks	Distributed anomaly detection
Watchdog and pathrater	A threshold-based scheme where nodes watch neighbors for signs of malicious activity. Once a threshold is crossed the malicious nodes are excluded from the network.	Distributed anomaly detection
Local intrusion detection system	Mobile agents use local SNMP data located in the management information base as audit sources for intrusion detection. Also implemented is the use of the Intrusion Detection Message Exchange Format (IDMEF) and a protocol for transporting such alerts, Intrusion Detection Exchange Protocol (IDXP), to ensure that intrusion detection systems running on a broad range of platforms can still interact and exchange intrusion-related information	Mobile-agent-based distributed anomaly detection
Intrusion detection architecture based on a stationary secure database	Secure, centralized, and stationary database used to store misuse signatures and user profiles. Uses mobile agents for intrusion detection	Mobile-agent-based compound detection
Distributed intrusion detection system based on mobile agents	Audit data from multiple sensors used to implement a bandwidth-conscious scheme for distributed intrusion detection using mobile agents.	Mobile-agent-based anomaly detection

■ **Table 2.** A summary of proposed IDSs.

the proposed system introduces new weaknesses in terms of excessive overheads for communication, storage, energy, or computation.

We see from Table 3 above that there is a trend to use mobile agents for intrusion detection and response in mobile ad hoc networks because these agents address the search and analysis problems involving multiple distributed resources in an efficient manner. As indicated by column 4, most proposed systems lack interoperability because they do not use the common message format that has been proposed by the IETF for communication between various IDS agents. Inter-IDS agent communication security is another area in which many of the proposed systems do not fare well.

Taking a leaf out of Axelsson's survey on IDSs [1] for wireline networks, we see the following dichotomies in system characteristics hold true for wireless ad hoc networks.

Time of detection: Two main groups can be identified in wireless ad hoc networks too: those that attempt to detect intrusions in real time and those that process audit data with some delay.

Locus of data processing: The audit data in general is processed, and new rules are derived from it in a distributed fashion. Each node in most of the surveyed systems takes the distributed approach to avoid being a single point of failure. The intrusion detection architecture is based on a secure stationary database being the only exception, where audit data is transferred to the stationary secure database with the help of mobile agents, and this audit data is then mined for new misuse patterns.

Security: The ability to withstand a hostile attack against the IDS itself. This area has been the subject of little investigation. With the trend toward using mobile agents for intrusion detection, most of the surveyed systems that use mobile agents still do not consider the security of the agent platform itself.

Degree of Interoperability: The degree to which the system can interoperate in conjunction with other IDSs, and accept audit data and reports from different sources. This is not the same as the number of different platforms on which the IDS itself runs. With the exception of one, most of the proposed systems are not interoperable with each other.

CONCLUSION

This article presents the current state of the art in the area of intrusion detection and response for wireless mobile ad hoc networks. Even though the research in intrusion detection started at least 15 years ago in the wired world, its application to wireless ad hoc networks is a rather recent development. Wireless ad hoc networks are intrinsically resource-constrained, which makes several of the schemes proposed in the wired world inadequate, as discussed earlier. Approaches that require analysis of large trace data or attack signatures (used by misuse detection) or centralized analysis engines are not preferable. Instead, schemes that are distributed and collaborative (e.g., anomaly-detection-based schemes) are likely to be more applicable. One key advantage of using an anomaly detection

Attributes of an ideal intrusion detection system for MANETs					
Proposed system↓	Introduces weaknesses/overheads	Fault-tolerant	Interoperates with other IDSs	Scalable	Detects new attack patterns
Distributed intrusion detection system for ad hoc networks	Yes. The use of encryption for communication between IDS agents will slow down the communication process.	Yes	No	Yes	Yes, but is ineffective against IP spoofing
Intrusion detection and response model for the AODV protocol	Yes. The use of encryption for communication between IDS agents will slow down the communication process	Yes	No	Yes	Yes
Techniques for Intrusion Resistant Ad Hoc Routing Algorithms	No	Yes	No	Yes	Yes
Watchdog and pathrater	No	Yes	No	Yes	Yes, but fails against collaborative attacks
Local intrusion detection system	No	No, security of the agent platform itself is not addressed.	Yes	Yes	Yes
Intrusion detection architecture based on a stationary secure database	Yes, limits the communication overhead but might increase the storage overhead in terms of misuse signatures on a node	No, the stationary static database could be a bottleneck/single point of failure	No	Yes	Yes, but with a delay because the stationary database has to mine the limited audit data given to it by agents and has to broadcast new signatures
Distributed intrusion detection system based on mobile agents	No	Yes	No	Yes	Yes

■ **Table 3.** Comparison of different proposed architectures against ideal characteristics for IDSs in MANETs

scheme is that it requires less modification of current routing protocols, and allows trace analysis and anomaly detection to be performed locally in each node. At present the IETF has a working group on intrusion detection [2] that covers current and future research topics, and is an excellent source for information on the scope of future work.

ACKNOWLEDGMENTS

The authors would like to thank the editors of this issue for allowing us to revise the article as per the reviewers' comments and working closely with us in improving the quality of the article. We would also like to thank the anonymous reviewers for their constructive criticism and suggestions on the article.

REFERENCES

- [1] S. Axelsson, "Intrusion Detection Systems: A Taxonomy and Survey," Tech. report no. 99-15, Dept. of Comp. Eng., Chalmers Univ. of Technology, Sweden, Mar. 20, 2003.
- [2] <http://www.ietf.org/html.charters/idwg-charter.html>
- [3] Y. Zhang and W. Lee, "Intrusion Detection in Wireless Ad Hoc Networks," *6th Int'l. Conf. Mobile Comp. and Net.*, Aug. 2000, pp. 275–83.
- [4] S. Bhargava and D. P. Agrawal, "Security Enhancements in AODV Protocol for Wireless Ad Hoc Networks," *VTC 2001 Fall*, vol. 4, Oct. 7–11, 2001, pp. 2143–47.
- [5] C. E. Perkins, E. M. Royer, and Samir R. Das, "Ad Hoc On-Demand Distance Vector Routing," IETF draft, Oct. 1999.
- [6] R. Ramanujan et al., "Techniques for Intrusion-Resistant Ad Hoc Routing Algorithms (TIARA)" *MILCOM 2000*, vol. 2, Oct. 22–25, 2000, pp. 660–64.

- [7] S. Marti et al., "Mitigating Routing Misbehavior in Mobile Ad Hoc Networks," *Proc. 6th Annual Int'l. Conf. Mobile Comp. and Net.*, Boston, MA, pp. 255–65.
- [8] D. B. Johnson and D. A. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks," *Mobile Computing*, T. Imielinski and H. Korth, Eds., Kluwer, 1996, pp. 153–81.
- [9] Y. Zhang, W. Lee, and Y.-A. Huang, "Intrusion Detection Techniques for Mobile Wireless Networks," *ACM J. Wireless Net.*, vol. 9, no. 5, Sept. 2003, pp. 545–56.
- [10] Y. Okazaki, I. Sato, and S. Goto, "A New Intrusion Detection Method based on Process Profiling," *Proc. SAINT 2002*, Jan. 28–Feb. 1, 2002, pp. 82–90.
- [11] P. Albers et al., "Security in Ad Hoc Networks: a General Intrusion Detection Architecture Enhancing Trust Based Approaches," *1st Int'l. Wksp. Wireless Info. Sys.*, Ciudad Real, Spain, Apr. 3–6, 2002.
- [12] A. B. Smith, "An Examination of an Intrusion Detection Architecture for Wireless Ad Hoc Networks," *5th Nat'l. Colloq. for Info. Sys. Sec. Education*, May 2001.
- [13] O. Kachirski and R. Guha, "Intrusion Detection Using Mobile Agents in Wireless Ad Hoc Networks," *Knowledge Media Net.*, *Proc. IEEE Wksp.*, July 10–12, 2002, pp. 153–58.

ADDITIONAL READING

- [1] C. Krugel and T. Toth, "Flexible, Mobile Agent Based Intrusion Detection for Dynamic Networks," *Euro. Wireless*, Italy, Feb. 2002.
- [2] S. Forrest et al., "A Sense of Self for Unix Processes," *Proc. 1996 IEEE Symp. Security and Privacy*, Los Alamitos, CA, 1996, pp. 120–28.
- [3] A. K. Ghosh and A. Schwartzbard, "A Study in using Neural Networks for Anomaly and Misuse Detection," *Proc. 8th USENIX Security Symp.*, 1999, pp. 141–52.
- [4] L. Zhou and Z. Haas, "Securing Ad Hoc Networks," *IEEE Network*, vol. 13, no. 6, Nov.–Dec. 1999, pp. 24–30.

One key advantage of using an anomaly detection scheme is that it requires less modification of current routing protocols and allows the trace analysis and anomaly detection to be performed locally in each node.

- [5] J. Kong *et al.*, Adaptive Security for Multi-layer Ad-Hoc Networks," Special Issue, *Wireless Commun. and Mobile Comp.*, 2002.
- [6] J.-P. Hubaux, L. Buttyan, and S. Capkun, "The Quest for Security in Mobile Ad Hoc Networks," *Proc. ACM Symp. Mobile Ad Hoc Net. and Comp.*, ACM Press, Oct. 2001, pp. 146–55.
- [7] E. H. Spafford and D. Zamboni, "Intrusion Detection Using Autonomous Agents," *Comp. Net.* 34, 2000, pp. 547–70.
- [8] W. Lee, S. J. Stolfo, and K. W. Mok, "A Data Mining Framework for Building Intrusion Detection Models," *Security and Privacy, Proc. 1999 IEEE Symp.*, May 9–12, 1999, pp. 120–32.

BIOGRAPHIES

AMITABH MISHRA [SM] (Mishra@vt.edu) received an M.Eng. degree in 1982 and a Ph.D. degree in 1985 in electrical engineering from McGill University, Montreal, Canada, and an M.S. degree in computer science in 1996 from the University of Illinois at Urbana-Champaign. He is currently an associate professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech, Blacksburg, where he conducts research primarily in wireless ad hoc and sensor networks, and mobile computing. He is currently investigating scalability issues in wireless sensor networks, energy-efficient MAC as well as cross-layer protocols for wireless sensor networks, intrusion detection and prevention algorithms for ad hoc networks, performance, interworking, and architectures of heterogeneous wireless networks, and quality of service issues for 3G/4G cellular networks. His other interests include wireless networking solutions for national critical infrastructure protection, fault tolerance, and software performance engineering. He has authored over 50 papers in journals and conferences, and has seven patents granted and/or currently pending. From 1987 to 2000 he was a member of technical staff at Lucent Technologies-Bell Laboratories, Naperville, Illinois, where his focus was on application architecture performance. AIN, ISDN, ATM, GPRS, CDMA2000, and UMTS were the major

areas he worked on while with Bell Laboratories. He received his B. Eng. and M.Tech. degrees in electrical engineering from Jabalpur University and Indian Institute of Technology, Kharagpur in 1973 and 1975, respectively. He is Chair of IEEE Communications Software and Secretary of IEEE Internet Technical Committees, and a member of ACM and SIAM. He has served on the program committees of at least 20 conferences over the past three years.

KETAN M. NADKARNI (ketann@vt.edu) received his B.Eng. degree in electronics and telecommunication engineering in 1999 from Goa University, India, and his M.S. degree in computer science and applications in 2003 from Virginia Tech. His research interests include wireless ad hoc networks and network security, and his Master's thesis focused on intrusion detection for wireless mobile ad hoc networks. His research accomplishments include a book chapter and a conference publication. From 1999 to 2001 he worked as a systems engineer at Tata Infotech Ltd., Bombay, India where his responsibilities included, but were not limited to, software development, consulting, and implementation. He is currently working as a software developer for Meridium Inc., Virginia, that specializes in enterprise asset performance management software for the process industry.

ANIMESH PATCHA (apatcha@vt.edu) received his B.Eng. degree in electrical and electronics engineering in 1998 from Birla Institute of Technology, India, and his M.S. degree in computer engineering from Illinois Institute Of Technology, Chicago, in 2002. From January 1999 to December 2000 he worked as a senior software developer at Zensar Technologies. He is currently a Ph.D. student in computer engineering at the Bradley Department of Electrical and Computer Engineering at Virginia Tech. His current research focus is on the application of mobile agents for intrusion detection. His main research interests are in the area of wireless network security and intrusion detection in mobile ad hoc networks.