Ad Hoc and P2P Security

Dempster-Shafer Theory for Intrusion Detection in Ad Hoc Networks

Without a fixed security infrastructure, mobile ad hoc networks must distribute intrusion detection among their nodes. But even though a distributed intrusion-detection system can combine data from multiple nodes to estimate the likelihood of an intrusion, the observing nodes might not be reliable. The Dempster-Shafer theory of evidence is well suited for this type of problem because it reflects uncertainty. Moreover, Dempster's rule for combination gives a numerical procedure for fusing together multiple pieces of evidence from unreliable observers. The authors review the Dempster-Shafer theory in the context of distributed intrusion detection and demonstrate the theory's usefulness.

Thomas M. Chen and Varadharajan Venkataramanan
Southern Methodist University

The nature of mobile ad hoc networks makes them vulnerable to a variety of attacks and difficult to protect.[1,2] Wireless links, for example, are susceptible to passive eavesdropping and active interference; the nodes' nomadic nature increases the risk of physical attack and compromise; and ad hoc networks — by definition — lack a fixed infrastructure to support security, which means security mechanisms must be implemented in the nodes themselves. Moreover, the need for security in ad hoc networks is heightened by nodal interdependencies. Attacks on nodes can disrupt communications, for example, but they're particularly worrisome if compromised nodes — those taken over by an intruder without any obvious sign of attack — can present the proper credentials and then interfere with secure routing protocols.[2,3]

Intrusion detection is important in any security framework, but implementing it in an ad hoc network is difficult due to the absence of any natural concentration points at which to monitor the network's traffic. Any ad hoc node can observe part of the total traffic, making distributed intrusion detection the most logical approach: each node is responsible for detecting intrusions in its local neighborhood.[4] In this form of intrusion detection, a node's neighbors observe that node's external behavior and form a judgment about the node's "trustworthiness." (We focus here on intrusion as the

1089-7801/05/$20.00 © 2005 IEEE
Published by the IEEE Computer Society

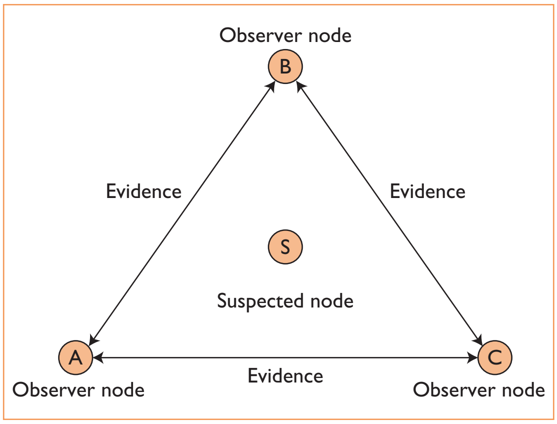

Figure 1. Trustworthiness scenario. The three nodes observing node S and combining their evidence could be untrustworthy themselves.

reason why a node behaves suspiciously, but such behavior might also be due to faults or malfunctions. Previously proposed intrusion-detection methods can't determine the behavior's true cause — compromise or fault — so researchers are also pursuing an alternative approach called *intrusion tolerance*, which seeks to maintain proper network operation in the face of hostile attacks.[5])

A common problem in distributed intrusion detection is how to combine observational data from multiple nodes that can vary in their reliability or trustworthiness. Previous approaches have used simplistic combination techniques such as averaging or majority voting (see the "Related Work in Intrusion Detection" sidebar). In this article, we investigate the Dempster-Shafer evidence theory, which is well suited to this type of problem for two reasons: first, it reflects uncertainty or a lack of complete information, and second, Dempster's rule for combination gives a convenient numerical procedure for fusing together multiple pieces of data.

Combining Evidence from Multiple Observers

Figure 1 shows an example of a trustworthiness situation among nodes. In this figure, nodes A, B, and C share their independent observations about suspected node S's behavior. We refer to the observation data as *evidence*, which can be in the form of malcounts (the number of observed occurrences of misbehavior) or some other rating. The distributed intrusion-detection system must somehow combine this evidence into a decision about node S's trustworthiness, but any of the three observing nodes could be untrustworthy themselves (due to compromise or some other reason) and thus contribute unreliable evidence. Note that a trustworthy node can also contribute unreliable evidence due to inaccurate intrusion detection; detection accuracy is an issue for any intrusion-detection system.

We can imagine several possible ways to combine evidence from multiple observers. We could simply average the evidence,[6] for example, but averaging ignores the fact that some observers are more trustworthy or reliable than others. Another simple method is a majority-decision rule (or majority voting),[7] which operates under the assumption that most of the nodes observing a suspect node are trustworthy. During evidence collection, any misbehaving observer could choose either to not offer any evidence or to provide misinformation in the form of falsified evidence; in either case, this minority of misbehaving observers wouldn't change the consensual decision reached by a majority of trustworthy observers.

Unfortunately, the restrictive assumptions for majority consensus might be infeasible to guarantee in practice — in a general setting, for example, it's hard to tell which observers are compromised. The Dempster-Shafer theory of evidence, originated by Arthur Dempster[8] and later revised by Glenn Shafer,[9] addresses this situation by representing uncertainty in the form of belief functions. The essential idea is that an observer can obtain degrees of belief about a proposition from a related proposition's subjective probabilities. The theory's practical appeal is due largely to Dempster's rule for combining beliefs based on independent pieces of evidence. Although extensive literature surrounds many applications of the Dempster-Shafer theory, very little appears to be applied to distributed intrusion detection except for sensor fusion.[10]

Bayesian Inference

The Dempster-Shafer theory is often illustrated in comparison with the better-understood Bayesian approach.[11] Bayesian inference has wide appeal because it's well grounded in the formalities of probability — namely, the well-known Bayes' theorem:

$$P(H \mid E) = \frac{P(E \mid H)P(H)}{P(E)}. \tag{1}$$

Bayesians interpret the a posteriori probability $P(H|E)$ as a measure of belief about a hypothesis or proposition H updated in response to evidence E. The a priori probability $P(H)$ reflects the belief about H in the absence of evidence.

Applied to Figure 1, suppose that nodes A, B,

Related Work in Intrusion Detection

We can trace back the ideas underlying distributed intrusion detection in mobile ad hoc networks to Watchers (Watching for Anomalies in Transit Conservation: a Heuristic for Ensuring Router Security).[1] The Watchers scheme was an early proposal for detecting misbehaving routers by using distributed network monitoring; specifically, the approach depends on each router monitoring the traffic passing through its neighboring routers. Although Watchers wasn't specifically intended for ad hoc networks, all nodes in ad hoc networks function as routers, so the Watchers approach is easily applicable.

A router can detect packets dropped by neighboring routers by comparing the observed amounts of traffic flowing into and out of a neighbor. Each router also counts any packets misrouted by neighboring routers, assuming that each router knows its neighbors' routing tables from a link-state routing protocol. The routers periodically share their respective data via a flooding protocol before starting a diagnostic phase. In this diagnostic phase, every router compares the counts collected from their neighboring routers to determine if any routers have

- misrouted too many packets,
- not participated correctly in the Watchers scheme,
- broadcasted counts that have discrepancies with their neighbors' counts, or
- appeared to drop more packets than a given threshold.

In response to any routers deemed to be misbehaving, their neighbors will change their routing tables to avoid forwarding packets through those misbehaving ones.

Yongguang Zhang and Wenke Lee's scheme has influenced several researchers.[2] In it, each node concurrently runs a software agent that monitors its own system activities as well as traffic among neighboring nodes within its radio range. Each node also analyzes its own data for local intrusion detection. Here, intrusion detection is based on statistical anomaly detection, rather than misuse detection, because of the perceived difficulties of continually updating misuse detection rules (or signatures) in an ad hoc network. If an intrusion warrants a broader investigation, nodes are expected to trigger the cooperation of other nodes for global-scale intrusion detection. The algorithm for performing this task collects observed data from all the nodes about the suspected node, then weighs the majority consensus to determine whether an intrusion has occurred.

Sergio Marti and colleagues proposed the idea of ad hoc nodes monitoring their

continued on p. 38

and C offer respective pieces of evidence e_A, e_B, and e_C to support the hypothesis that node S is trustworthy. The a posteriori probability would be

$$P(H \mid e_A, e_B, e_C) = \frac{P(e_A, e_B, e_C \mid H)P(H)}{P(e_A, e_B, e_C \mid H)P(H) + P(e_A, e_B, e_C \mid \bar{H})(1 - P(H))}, \qquad (2)$$

where $\bar{H}$ is the hypothesis "not H" (that is, S is untrustworthy). Researchers often assume that the observer nodes are conditionally independent of each other, meaning that they make independent observations of the same fact. The computation of Equation 2 is then simplified by the factorization $P(e_A, e_B, e_C|H) = P(e_A|H)P(e_B|H)P(e_C|H)$.

Clearly, the Bayesian approach requires complete knowledge of both prior and conditional probabilities, which might be difficult to determine in practice. We often estimate prior probabilities from empirical data, or, in the absence of empirical data, we assume them to be uniform or some other distribution. The outcome reflects these assumptions, so the Bayesian approach's critics often point out that the method isn't well equipped to handle states of ignorance.

Dempster-Shafer Formalities

The Dempster-Shafer theory is appealing partly because it can handle uncertainty or ignorance — that is, the lack of knowledge of the complete probabilistic model required for Bayesian inference. As an introduction to reasoning in Dempster-Shafer, suppose that node A is either trustworthy with probability 0.8 or untrustworthy with probability 0.2. Also, suppose that node A states that suspected node S is trustworthy. If node A itself is trustworthy, then its claim is accurate, but if A isn't trustworthy, its claim isn't necessarily inaccurate. Node A's claim gives evidence for 0.8 degrees of belief in node S's trustworthiness, but a zero degree of belief (not 0.2) that node S is untrustworthy. The zero doesn't imply an impossibility (as in a zero probability); it means that node A's evidence gives no support to the belief that node S is untrustworthy. An interval bounded by 0.8 and zero might constitute a type of belief function. In light of uncertainty, Dempster-Shafer is concerned with bounds for probabilities of provability rather than computing probabilities of truth. The two bounds used in Dempster-Shafer are called *belief* and *plausibility*, as we'll describe in more detail later.

A *frame of discernment* (also called a *universe*

Related Work in Intrusion Detection (cont.)

continued from p. 37

neighboring nodes' packet-forwarding behavior in what they called a *watchdog* process.[3] After a node forwards a packet, the watchdog monitors the next node to see that the packet is forwarded again. The scheme assumes source routing, with each packet carrying its route information so that the watchdog knows a tracked packet's proper route. If a watchdog sees a neighboring node drop more packets than a given threshold, the node is deemed to be misbehaving. Because a watchdog is a rather simple monitoring process, Marti and his colleagues found several limitations with it.

Sonali Bhargava and Dharma Agrawal's system is essentially an enhancement of Zhang and Lee's approach.[4] Each node maintains a "malcount" for neighboring nodes, which is the number of observed occurrences of misbehavior. When a node's malcount exceeds a given threshold, their neighbors send out an alert to the other nodes, which then check their malcounts for the suspected node (this can, in turn, trigger secondary alerts). If a suspected node triggers two or more alerts, it's deemed to be malicious. Naturally, this scheme works only if at least two trustworthy nodes observe a suspected node; it can fail if malicious nodes send out false alerts.

Sonja Buchegger and Jean-Yves LeBoudec proposed the Confidant (Cooperation of Nodes: Fairness in Dynamic Ad-hoc Networks) scheme, which, similar to previous approaches, relies on ad hoc nodes to monitor their neighboring nodes' routing behavior.[5] Source routing is assumed, so nodes know the correct route for tracked packets, but Confidant's innovation is a reputation system that works with network monitoring and that consists of a table of observed nodes and their reputation ratings. If a node is observed to be misbehaving (deviating from expected routing behavior), the reputation system changes the node's rating by a weighting function depending on the new observation's trustworthiness.

Frank Kargl and colleagues' MobIDS (Mobile Intrusion Detection System) is generally similar to the other schemes described here.[6] Multiple sensors in the ad hoc network keep track of observed instances of the nodes' behavior, but in MobIDS, counts from multiple sensors are combined with a weighting function reflecting different sensors' credibility to create a local rating for a suspect node. These local ratings are then distributed periodically via broadcasting to the neighboring nodes. Each node averages the local ratings it receives into global ratings for other nodes, and nodes are deemed to be misbehaving if their ratings drop below a given threshold.

References

1. K. Bradley et al., "Detecting Disruptive Routers: A Distributed Network Monitoring Approach," *IEEE Network*, vol. 12, no. 5, 1998, pp. 50–60.
2. Y. Zhang and W. Lee, "Intrusion Detection in Wireless Ad-Hoc Networks," *Proc. 6th Ann. ACM Int'l Conf. Mobile Computing and Networking*, ACM Press, 2000, pp. 275–283.
3. S. Marti et al., "Mitigating Routing Misbehavior in Mobile Ad Hoc Networks," *Proc. 6th Ann. ACM Int'l Conf. Mobile Computing and Networking*, ACM Press, 2000, pp. 255–265.
4. S. Bhargava and D. Agrawal, "Security Enhancements in AODV Protocol for Wireless Ad Hoc Networks," *Proc. IEEE Vehicular Tech. Conf.*, IEEE CS Press, 2001, pp. 2143–2147.
5. S. Buchegger and J-Y. Le Boudec, "Performance Analysis of the CONFIDANT Protocol (Cooperation of Nodes: Fairness in Dynamic Ad-Hoc Networks)," *Proc. 3rd ACM Int'l Symp. Mobile Ad Hoc Networking and Computing*, ACM Press, 2002, pp. 226–236.
6. F. Kargl et al., "Sensors for Detection of Misbehaving Nodes in MANETs," *Proc. Detection of Intrusion and Malware and Vulnerability Assessment* (DIMVA 2004), 2004; www.dimva.org/dimva2004/.

of discourse) in Dempster-Shafer is a set of mutually exclusive and exhaustive possibilities denoted by Ω, which is similar to a state space in probability. Any hypothesis A will refer to a subset of Ω for which observers can present evidence. The set of all possible subsets of Ω, including itself and the null set $\varnothing$, is called a *power set* and designated as 2^Ω. Thus, the power set consists of all possible hypotheses or so-called focal elements $2^\Omega = \{A_1, \Omega, A_n\}$.

We can assign hypotheses to any of three types of values. Basic probability numbers (also called basic belief mass) are a mapping of each hypothesis A to a value $m(A)$ between 0 and 1, such that

- the basic probability number of the null set $\varnothing$ is $m(\varnothing) = 0$, and
- the sum $m(A_1) + \ldots + m(A_n) = 1$.

We can interpret the basic probability number $m(A)$ as the portion of total belief assigned to hypothesis A, reflecting the evidence's strength of support.

The second type of assignment is a belief function that maps each hypothesis B to a value $bel(B)$ between 0 and 1, defined as

$$bel(B) = \sum_{j:A_j \subset B} m(A_j). \tag{3}$$

The belief function represents the weight of evidence supporting B's provability.

The third type of assignment is a plausibility function that maps each hypothesis B to a value $pls(B)$ between 0 and 1, defined as

$$pls(B) = \sum_{j:A_j \cap B \neq \varnothing} m(A_j). \tag{4}$$

The plausibility function is the weight of evidence that doesn't refute B, and belief and plausibility are related by

$$pls(B) = 1 - bel(\bar{B}), \tag{5}$$

where $\bar{B}$ is the hypothesis "not B." Shafer showed that a one-to-one correspondence exists between basic probability numbers, belief, and plausibility, meaning that any of the three functions is sufficient for deriving the other two.

Dempster's Rule for Combination

Dempster's rule for combination is a procedure for combining independent pieces of evidence. Let's revisit our earlier example in which node A is trustworthy with probability 0.8 and suspected node S is trustworthy. After collecting evidence from node A, suppose that node B gives its own evidence about node S. Assume node B is either trustworthy with probability 0.8 or untrustworthy with probability 0.2, independently of node A. If A and B both claim that S is trustworthy, then the hypothesis that node S is trustworthy will be true if at least one of the observers (A or B) is trustworthy. The probability that at least one observer is trustworthy is 1 – (0.2)(0.2) = 0.96, thus the degree of belief would be 0.96 in S's trustworthiness in this case.

On the other hand, suppose that node A claims that S is trustworthy whereas node B claims that it isn't. Nodes A and B can't both be correct, thus they can't both be trustworthy — either one is trustworthy or neither is. Given the prior probabilities of their trustworthiness, let's calculate the posterior probabilities given that both can't be trustworthy. In this case, the posterior probability that only A is trustworthy is

$$\frac{(0.8)(0.2)}{1-(0.8)(0.8)} = \frac{4}{9}.$$

Likewise, the posterior probability that only B is trustworthy is 4/9. The probability that neither A nor B is trustworthy is 1/9, thus the degree of belief would be 4/9 that node S is trustworthy (believing node A) and 4/9 that it isn't (believing node B).

More formally, suppose $m_1(A)$ and $m_2(A)$ are the basic probability numbers from two independent observers (in the same frame of discernment). Dempster's rule for combination consists of the orthogonal sum

$$m(B) = m_1(B) \oplus m_2(B) = \frac{\Sigma_{i,j:A_i \cap A_j = B}\, m_1(A_i) m_2(A_j)}{\Sigma_{i,j:A_i \cap A_j = \varnothing}\, m_1(A_i) m_2(A_j)}. \tag{6}$$

We can combine more than two belief functions pairwise in any order.

Dempster-Shafer Applied to Distributed Intrusion Detection

For a simple illustration of Dempster-Shafer, suppose that the frame of discernment consists of two possibilities concerning suspected node S: $\Omega = \{\mathrm{T}, \bar{T}\}$, where T means node S is trustworthy, and $\bar{T}$ means it isn't. For this Ω, the power set has three focal elements: hypothesis $H = \{T\}$ that S is trustworthy; hypothesis $\bar{H} = \{\bar{T}\}$ that it isn't; and (universe) hypothesis $U = \Omega$ that S is either trustworthy or untrustworthy. Suppose the probability of node A being trustworthy is α. If node A claims that S is trustworthy, then its basic probability assignment will be

$$\begin{aligned} m_1(H) &= \alpha \\ m_1(\bar{H}) &= 0 \\ m_1(U) &= 1 - \alpha. \end{aligned} \tag{7}$$

If node A claims that S is untrustworthy, its basic probability assignment will be

$$\begin{aligned} m_1(H) &= 0 \\ m_1(\bar{H}) &= \alpha \\ m_1(U) &= 1 - \alpha. \end{aligned} \tag{8}$$

Likewise, given prior probabilities for the trustworthiness of nodes B and C, we would construct their basic probability assignments m_2 and m_3 similarly.

Next, the combined belief of A, B, and C in H is $bel\,(H) = m(H) = m_1(H) \oplus m_2(H) \oplus m_3(H)$ following Dempster's rule for combination (Equation 6). We can compute this by combining any pair of arguments and then combining the result with the remaining third argument. For example, let's first combine m_1 and m_2:

$$m_1(H) \oplus m_2(H) = \frac{1}{K}[m_1(H)m_2(H) + m_1(H)m_2(U) + m_1(U)m_2(H)]$$

$$m_1(\bar{H}) \oplus m_2(\bar{H}) = \frac{1}{K}[m_1(\bar{H})m_2(\bar{H}) + m_1(\bar{H})m_2(U) + m_1(U)m_2(\bar{H})]$$

$$m_1(U) \oplus m_2(U) = \frac{1}{K} m_1(U) m_2(U), \tag{9}$$

where

$$K = m_1(H)m_2(H) + m_1(H)m_2(U) + m_1(U)m_2(H) + m_1(\bar{H})m_2(\bar{H}) + m_1(\bar{H})m_2(U) + m_1(U)m_2(\bar{H}) + m_1(U)m_2(U). \tag{10}$$

We can similarly combine the result from Equation 9 with m_3.

To weigh and combine *A*, *B*, and *C*'s statements about hypothesis *H*, the Dempster-Shafer approach must know *A*, *B*, and *C*'s trustworthiness or reliability. Contrast this with the Bayesian approach in Equation 2 — to combine the evidence offered about *H*, we need to know the prior probability *P*(*H*) as well as every conditional probability that observer *i* would offer evidence e_i when *H* is true and when it isn't. Compared to Dempster-Shafer, the Bayesian approach requires much more information, which might not even be available.

An attempt to use the Bayesian approach to form a judgment about *S* is problematic without additional information or assumptions.

An Example

The previous section described how to apply Dempster-Shafer to distributed intrusion detection, but we haven't clearly shown how it works in practice or why it's better than simple majority voting. Again, let's consider the situation in Figure 1, in which a sensor (located perhaps in another local node) wants to combine evidence from nodes *A*, *B*, and *C* about a suspected node *S*. In each case, *A*, *B*, and *C* could be trustworthy or untrustworthy. We can compute an initial estimate of the nodes' trustworthiness by keeping a malcount for each of them and then comparing the malcounts to a set of thresholds; a malcount exceeding higher thresholds lowers that node's trustworthiness rating. We can adjust their trustworthiness ratings later by applying the same procedure for combining evidence about *S* and then using that procedure to judge each observer (*A*, *B*, and *C*) in turn. Another way to adjust a node's trustworthiness rating is to look for cases in which it offered evidence contrary to the eventual judgment (for example, if *S* was judged to be trustworthy, but an observer stated that it wasn't).

Although we can use a simple majority-decision rule for combining evidence — if at least two observers state that *S* is trustworthy or untrustworthy, the judgment about *S* will follow the majority — a correct majority decision requires a majority of observers to offer accurate evidence. The Dempster-Shafer approach doesn't have this limitation. Moreover, majority voting's final judgment is a simple binary decision about whether *S* is trustworthy. Dempster-Shafer produces a judgment value between 0 and 1 that reflects the degree of belief in that judgment.

An obvious alternative to majority voting might be averaging the observers' numerical evidence — each observer could offer a number between 0 and 1 to vote on *S*'s trustworthiness, with the consensus judgment being the average of those numbers. However, compromised observers could cause an error in the final judgment by offering deliberately incorrect votes. Dempster-Shafer's advantage here is that it discounts evidence from untrustworthy or uncertain observers.

As we mentioned earlier, an attempt to use the Bayesian approach to form a judgment about *S* is problematic without additional information or assumptions. One difficulty is the requirement to know the a priori probability that *S* is trustworthy or the belief that *S* is trustworthy in the absence of any evidence. Without evidence, we might assume that *S* is equally likely to be trustworthy or not, but this assumption could lead to a completely erroneous result. Another difficulty is the requirement of knowing the conditional probabilities that an observer *i* will offer evidence e_i given that *S* is trustworthy or not. It's unclear how to determine these conditional probabilities.

Case One

Suppose that the ratings for nodes *A*, *B*, and *C* indicate that they're trustworthy with probabilities 0.9, 0.8, and 0.2, respectively. Also, suppose that *S* is actually trustworthy. If nodes *A* and *B* agree that suspect node *S* is trustworthy whereas node *C* disagrees, their combined degree of belief in *S*'s trustworthiness turns out to be high, at 0.975, but node *C*'s evidence is discounted in Dempster-Shafer as having a high amount of uncertainty.

In this case, simple majority voting works as well because most of the observers are trustworthy (and assumed to offer accurate evidence).

Case Two

Suppose now that both nodes B and C are trustworthy with probability 0.2, whereas node A is still trustworthy with probability 0.9. If node A claims that S is trustworthy, whereas nodes B and C claim that it isn't, the combined degree of belief in S's trustworthiness turns out to still be high at 0.878. Although the untrustworthy nodes B and C form a majority, their evidence is substantially discounted in the combined belief function. In this case, simple majority voting would lead to an incorrect judgment.

Case Three

The previous cases show that if the observing nodes' trustworthiness ratings are accurate, Dempster-Shafer will properly discount the evidence from untrustworthy nodes. But what happens if all the observers are trusted, yet their evidence disagrees because one node is inaccurate? Suppose all observing nodes are trustworthy with probability 0.8. Nodes A and B claim that S is trustworthy whereas node C disagrees. The combined degree of belief in S's trustworthiness turns out to still be high at 0.828. In this case, the evidence is weighed equally, and the final judgment essentially follows the majority consensus. This case shows that Dempster-Shafer is tolerant of trusted but inaccurate evidence as long as most of the evidence is accurate.

Some people claim that Dempster-Shafer is an extension or generalization of Bayesian theory. Although this claim's validity is debatable, Dempster-Shafer does seem to offer some practical advantages. It offers a mathematical way to combine evidence from multiple observers without the need to know about a priori or conditional probabilities as in the Bayesian approach.

One area of difficulty that we continue to study is the problem of determining initial estimates of nodes' trustworthiness. Dempster-Shafer can combine observations from trustworthy and untrustworthy nodes, but the results depend on accurate initial estimations of each observer's trustworthiness. Another problem is how to ensure that the judgments about each node's trustworthiness, formed from their neighbors' observations, will be consistent and resistant to sabotage efforts by malicious nodes.

References

1. L. Zhou and Z. Haas, "Securing Ad Hoc Networks," *IEEE Network*, vol. 13, no. 6, 1999, pp. 24–30.
2. H. Deng, W. Li, and D. Agrawal, "Routing Security in Wireless Ad Hoc Networks," *IEEE Comm. Magazine*, vol. 40, no. 10, 2002, pp. 70–75.
3. K. Sanzgiri et al., "Authenticated Routing for Ad Hoc Networks," *IEEE J. Selected Areas in Comm.*, vol. 23, no. 3, 2005, pp. 598–610.
4. A. Mishra, K. Nadkarni, and A. Patcha, "Intrusion Detection in Wireless Ad Hoc Networks," *IEEE Wireless Comm.*, vol. 11, no. 2, 2004, pp. 48–60.
5. C. Basile, Z. Kalbarczyk, and R. Iyer, "Neutralization of Errors and Attacks in Wireless Ad Hoc Networks," *Proc. Int'l Conf. Dependable Systems and Networks* (DSN 2005), IEEE CS Press, 2005, pp. 518–527.
6. F. Kargl et al., "Sensors for Detection of Misbehaving Nodes in MANETs," *Proc. Detection of Intrusion and Malware and Vulnerability Assessment* (DIMVA 2004), 2004; www.dimva.org/dimva2004/.
7. Y. Zhang and W. Lee, "Intrusion Detection in Wireless Ad-Hoc Networks," *Proc. 6th Ann. ACM Int'l Conf. Mobile Computing and Networking*, ACM Press, 2000, pp. 275–283.
8. A. Dempster, "Upper and Lower Probabilities Induced by a Multivalued Mapping," *Ann. Mathematical Statistics*, vol. 38, no. 2, 1967, pp. 325–339.
9. G. Shafer, *A Mathematical Theory of Evidence*, Princeton Univ. Press, 1976.
10. H. Wu et al., "Sensor Fusion Using Dempster-Shafer Theory," *Proc. IEEE Instrumentation and Measurement Conf.*, IEEE CS Press, 2002, pp. 7–12.
11. J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 1988.

Thomas M. Chen is an associate professor in the Department of Electrical Engineering at Southern Methodist University. His research is in network security and traffic control. Chen has a PhD in electrical engineering from the University of California, Berkeley. He's also the associate editor in chief of *IEEE Communications Magazine*, a senior technical editor for *IEEE Network*, and a past associate editor for *ACM Transactions on Internet Technology*. Contact him at tchen@engr.smu.edu.

Varadharajan Venkataramanan is a PhD student in the School of Engineering at Southern Methodist University. His research interests are in security issues in ad hoc networks, quality of service in wireless and wired networks, and stochastic processes. Venkataramanan has a B.Eng. in electronics and telecommunication from the University of Mumbai, India, and an MS in electrical engineering from SMU. Contact him at venvar@engr.smu.edu.