Self-Policing Mobile Ad-hoc Networks

Sonja Buchegger and Jean-Yves Le Boudec

EPFL-IC-LCA

CH-1015 Lausanne, Switzerland

sonja.buchegger@epfl.ch, jean-yves.leboudec@epfl.ch

Abstract:

Misbehavior in mobile ad-hoc networks occurs for several reasons. Selfish nodes misbehave to save power or to improve their access to service relative to others. Malicious intentions result in misbehavior as exemplified by denial of service attacks. Faulty nodes simply misbehave accidentally. Regardless of the motivation for misbehavior its impact on the mobile ad-hoc network proves to be detrimental, decreasing the performance and the fairness of the network, and in the extreme case, resulting in a non-functional network. Countermeasures to prevent or to combat misbehavior have been proposed, such as payment schemes for network services, secure routing protocols, intrusion detection and reputation systems to detect and isolate misbehaved nodes. We discuss the trade-offs and issues of self-policing mobile ad-hoc networks and give an overview of the state of the art, discussing and contrasting several solution proposals.

1 Introduction and Chapter Overview

In mobile ad-hoc networks, nodes act as both routers and terminals. For the lack of routing infrastructure, they have to cooperate to communicate. Cooperation at the network layer takes place at the level of routing, i.e. finding a path for a packet, and forwarding, i.e. relaying packets for other nodes.

Misbehavior means aberration from regular routing and forwarding behavior resulting in detrimental effects on the network performance. Misbehavior arises for several reasons. When a node is faulty its erratic behavior can deviate from the protocol and thus produce non intentional misbehavior. Intentional misbehavior aims at providing an advantage for the misbehaved node. An example for an advantage gained by misbehavior is power saved when a selfish node does not forward packets for other nodes. An advantage for a malicious node arises when misbehavior enables it to mount an attack.

The detrimental effects of misbehavior result in unfairness and degraded performance and they can endanger the functioning of the entire network. In order to avoid these adverse effects, mobile ad-hoc network routing protocols have to be able to cope with misbehavior attempts. Section 2 discusses types of misbehavior, their payoff for the attacker, and their effect on the network. It thus gives a motivation for the need of enhancements of mobile ad-hoc networks to counteract misbehavior.

In the literature, three main approaches to strengthen mobile ad-hoc networks have been proposed, namely payment systems, secure routing using cryptography, and detection and reputation systems for self-policing. Payment schemes serve as an incentive to forward packets for other nodes. Secure routing aims at the prevention of attacks by using cryptography to secure the routing messages them-

selves. Detection and reputation systems identify misbehaved nodes and isolate them from the network by monitoring and keeping records of past behavior. Section 3 gives an overview and comparison of these three solution tracks.

In Section 4 we discuss detection and reputation systems in more detail. They address the cases that have not been prevented. Not all types of misbehavior can be prevented, however, for misbehavior reputation systems it suffices that misbehavior can be detected. We present our approach to a self-policing mobile ad-hoc network and use it as a basis to compare with other proposed solutions.

There are several challenges to the design of misbehavior reputation systems, a fundamental example being that the system should not add vulnerabilities to the mobile ad-hoc routing protocol it is built to protect. There are trade-offs to take into account when considering which type of reputation to use, whether and how to use second-hand information, and so forth. We discuss the main challenges and issues in Section 4.2 and conclude in Section 5.

2 Node Misbehavior in Mobile Ad-hoc Networks

This section gives a more detailed problem description and reasons why it is a worthwhile question. Mobile ad-hoc networks have properties that render them more vulnerable to attacks and misuse, as we show in Section 2.1. Several attacks on routing and forwarding in mobile ad-hoc networks are described in Section 2.2. Finally, in Section 2.3, we illustrate the effects of misbehavior on the mobile ad-hoc network as well as the potential effects of countermeasures such as incentives for cooperation.

2.1 Reasons and Enablers for Misbehavior

The lack of infrastructure and organizational environment of mobile ad hoc networks offer special opportunities to intentionally misbehaved nodes. Without proper countermeasures, it is possible to gain various advantages by malicious behavior: better service than cooperating nodes, monetary benefits by exploiting incentive measures or trading confidential information; saving power by selfish behavior, preventing someone else from getting proper service, extracting data to get confidential information, and so on. Even if the misbehavior is not intentional, as in the case of a faulty node, the effects can be detrimental to the performance of a network.

Mobile ad-hoc networks have the following properties that can be exploited:

Lack of infrastructure. Nodes have to cooperate in the routing and forwarding of packets.

No organizational authorities. Any node can join an unmanaged mobile ad-hoc network, there is no access control and no specific entry point.

No central authorities. No permanent access to central services such as certification authorities can be assumed.

Wireless network. Nodes can promiscuously eavesdrop on communications by others. Collisions can be intentional or accidental.

Mobility. With high mobility routes are not valid over extended periods of time. Link errors can be ambiguous, communications can fail due to a node having moved out of range or due to an intentional interruption.

Routing protocols lack security. Most of the proposed routing protocols, such as DSR [12] and AODV [21], do not provide any security. Routing messages can be modified or fabricated, sent at inappropriate times or be omitted when needed. We discuss some proposals to add security in Section 3.2 and more detailed misbehavior descriptions in Section 2.2.

Potentially low battery power. Truly mobile and not merely portable devices have to be reasonably small and lightweight and therefore are often assumed to have limited battery power. This results in communications and computations being relatively expensive in power, opening the door to attacks aiming at excessive resource consumption of the target node, selfish behavior of resource conscious nodes, and limited ability to perform cryptographic computations.

In addition to authentication, integrity, confidentiality, availability, access control and non-repudiation (see [28] for details), which have to be addressed differently in a mobile, wireless, battery-powered and distributed environment, mobile ad hoc networks raise the issue of cooperation and fairness.

There is a trade-off between good citizenship, i.e. cooperation, and resource consumption, so nodes have to economize on their resources. Assuming rational behavior with a node maximizing its utility, the best strategy is not forward for other nodes. If several nodes, however, follow this strategy, the performance of the network deteriorates. In the extreme case of all nodes choosing this strategy, no communications can take place. This outcome is clearly unfavorable for the nodes. In game theoretic terms, this is a dilemma. Incentives are required to stimulate the cooperation among nodes.

2.2 Attacks

Ning and Sun [19] classify attacks on routing protocols as atomic attacks (modifying or forging one message) and compound attacks (combining or repeating several atomic attacks).

They argue that preventive security may not be enough to cope with insider attacks, where nodes can be compromised despite tamper-proof hardware.

They give a list of goals for an attacker, then look at the atomic attacks to see whether they can achieve them and also pick some compound attacks and investigate their effectiveness in reaching the four goals:

- route disruption
- route invasion
- node isolation
- resource consumption

Simulation results confirm the theoretical success of even atomic attacks, at least temporarily. For sustained success, e.g. to circumvent local route repair mechanisms, atomic attacks can be repeated.

Total non-cooperation with other nodes and only exploiting their readiness to cooperate is one of several boycotting behavior patterns. Therefore, there has to be an incentive for a node to forward messages that are not destined to itself. Attacks include incentive mechanism exploitation by message interception, copying, or forging; incorrect forwarding; and bogus routing advertisement.

Traffic diversion: Routes should be advertised and set up adhering to the chosen routing protocol and should truthfully reflect the knowledge of the topology

of the network. By diverting the traffic in the following ways, nodes can work against that requirement:

To get information necessary for successful malicious behavior, nodes can attract traffic to themselves or their colluding nodes by means of false routing advertisements. Although only suitable for devices that have enough power, a lot of information can be gathered this way by malicious nodes for later use to enable more sophisticated attacks.

Denial-of-service attacks can be achieved by bogus routing information (injecting of incorrect routing information or replay of old routing information or ‘black hole routes’) or by distorting routing information to partition the network or to load the network excessively, thus causing retransmissions.

Nodes can decide to forward messages to partners in collusion for analysis, disclosure, or monetary benefits; or may decide not to forward messages at all, thus boycotting communications.

In general, the following types of misbehavior can be indicated:

- no forwarding (of control messages nor data),
- unusual traffic attraction (advertises many very good routes or advertises routes very fast, so they are deemed good routes),
- deflecting traffic in order not to be used on a route,
- route salvaging (i.e. rerouting to avoid a broken link), although no error has been observed,
- lack of error messages, although an error has been observed,
- fabricating error messages, although no error has been observed,

- unusually frequent route updates,
- silent route change (tampering with the message header of either control or data packets).

Several more attacks have been proposed in the literature, such as the following:

Black hole. Reroute a path so that it ends up or passes a non existing node.

Grey hole. Like the black hole, but only performed sporadically.

Sleep deprivation. Make a node send messages excessively in order to decrease its resources.

2.3 The Effect of Misbehavior

Without appropriate countermeasures, the effects of misbehavior have been shown by several simulations [4, 17, 19] to be dramatically decrease network performance. Depending on the proportion of misbehaved nodes and their specific strategies network throughput can be severely degraded, packet loss increases, nodes can be denied service, and the network can be partitioned. Quantitative measures make more sense in comparison to routing protocols that have been enhanced with measures against misbehavior. We discuss these in Section 3.

In a theoretical analysis of how much cooperation mechanisms can help by increasing the probability of a successful forward, Lamparter, Plaggemeier, and Westhoff find that increased cooperation super-proportionally increases the performance for small networks (i.e. fairly short routes). Cooperation increases more if the initial probability e (the probability to cooperate by forwarding) is fairly

acceptable (above 0.6). Even small increases in e as given by δi, the change of the probability to cooperate in the presence of an incentive mechanism such as a reputation system, can have a dramatic improvement.

Zhang and Lee [30] argue that prevention measures, such as encryption and authentication, can be used in ad-hoc networks to reduce intrusions, but cannot eliminate them. Fore example, encryption and authentication cannot defend against compromised mobile nodes, which carry the private keys. No matter how many intrusion prevention measures are inserted in a network, there are always some weak links that one could exploit to break in. Intrusion detection presents a second wall of defense and it is a necessity in any high-survivability network.

3 Overview: Main Solution Tracks

The main solution tracks addressing the problem of misbehavior in mobile ad-hoc networks are secure routing, economic incentives, and detection and reputation systems. Economic incentives such as payment or counter schemes specifically address forwarding of packets for other nodes. Secure routing aims at securing the establishment and maintenance of routes.

Self-policing schemes aim at reactively detecting misbehavior and proactively isolating misbehaved nodes to prevent further damage. They are not restricted to any particular kind of misbehavior. The only requirement is that the misbehavior be detectable, i.e. observable and classifiable as such with a high probability.

In the following sections we describe the main features of some proposals within the respective solution tracks, briefly describe how they work, what they protect, and what the open problems are.

3.1 Payment Systems

Several approaches to provide economic incentives for cooperation have been proposed. They thus target the problem of selfish misbehavior. The main assumption is that nodes are economically rational.

Buttyán and Hubaux proposed incentives to cooperate by means of so-called **nuglets** [6] that serve as a per-hop payment in every packet in a secure module in each node to encourage forwarding. The secure module is required to ensure the correct number of nuglets is withdrawn or deposited. They propose two models for the payment of packet forwarding, the Packet Purse Model and the Packet Trade Model. In the Packet Purse Model the sender pays and thus loads the packet with a number of nuglets. Each intermediate node takes one nuglet when it forwards the packet. If there are no nuglets left at an intermediate node, the packet is dropped. If there are nuglets left in the packet once it reaches the destination, the nuglets are lost. In the Packet Trade Model, the destination pays for the packet. Each intermediate node buys a packet from the previous hop and sells it to the next for more nuglets. Since charging the destination and not the sender can lead to an overload of the network and the destination receiving packets it does not want, mainly the Packet Purse Model is considered. This model, however, can lead to the loss of nuglets which have to be re-introduced into the network by a central authority.

To address this problem, the authors introduced another approach based on credit **counters** [7], also implemented in tamper-proof hardware. In this approach, each node keeps track of its remaining battery power and credit. One of their findings of a simulation study of four different rules is that increased cooperation is beneficial not only for the entire network but also for individual nodes.

Zhong, Chen, and Yang proposed **Sprite** [31]. As opposed to nuglets or counters they do not require tamper-proof hardware to prevent the fabrication of payment units, but their payment scheme requires a central credit clearance service (CCS) to be available eventually. Nodes keep a receipt of a message when they receive it. The receipt contains a hash of the message itself so it can be verified which message the receipt belongs to. To claim their payment nodes have to send this receipt to the CCS. The CCS charges the sender based on the number of receipts, the number of intermediate nodes left to reach the destination, if any, and whether the destination has sent a receipt. The specific calculation of the fee is designed to make misbehavior in Sprite itself economically undesirable, even in the case of collusion. The sender then pays the nodes that sent a receipt to the CCS. For the nodes that were on the route but did not send a receipt, the sender has to pay a small fee to the CCS. In addition to the availability of a central authority, Sprite assumes source routing, and a public key infrastructure. They do not explain how the payment from the sender to nodes is done, whether nodes have accounts with the CCS which transfers the payment or whether nodes remunerate one another directly. In the latter case the money has to be unforgeable and payment has to be ensured.

Raghavan and Snoeren propose **priority forwarding** as incentives against selfish misbehavior. In their approach, potential dangers for ad-hoc networks are distinguished as misbehaving and greedy, where misbehavior constitutes a deviation from the protocol and should be taken care of by secure routing mechanisms. For greedy behavior, which is located at a higher layer in this approach, incentives to get priority forwarding are proposed to be given by payment.

3.2 Secure Routing with Cryptography

Secure routing proposals have been proposed mainly as modifications to existing routing protocols such as DSR [12] and AODV [21]. They aim at securing the routing messages by cryptographic means to prevent misbehavior by malicious nodes.

SRP, the Secure Routing Protocol by Papadimitratos and Haas [24], guarantees correct route discovery, so that fabricated, compromised, or replayed route replies are rejected or never reach the route requester. SRP assumes a security association between end-points of a path only, so intermediate nodes do not have to be trusted for the route discovery. This is achieved by requiring that the request along with a unique random query indentifier reach the destination, where a route reply is constructed and a message authentication code is computed over the path and returned to the source. The correctness of the protocol is proven analytically.

ARIADNE, a secure on-demand routing protocol by Hu, Perrig, and Johnson [11], prevents attackers from tampering with uncompromised routes consisting of uncompromised nodes. It is based on Dynamic Source Routing (DSR) and relies on symmetric cryptography only. It uses a key management protocol called TESLA that relies on synchronized clocks. Simulations have shown that the performance is close to DSR without optimizations.

SEAD, Secure Efficient Distance vector routing for mobile ad-hoc networks by Hu, Johnson and Perrig [10] is based on the design of destination-sequenced distance-vector routing (DSDV) and uses one-way hash functions to prevent uncoordinated attackers from creating incorrect routing state in another node. Performance evaluation has shown that SEAD outperforms DSDV-SQ in terms of packet delivery ratio, but SEAD adds overhead and latency to the network.

The **Security-aware Ad-hoc Routing (SAR)** protocol by Yi, Naldburg, and Kravets [29] modifies AODV to include security metrics for path computation and selection. They define trust levels according to organizational hierarchies with a shared key for each level, so that nodes can state their security requirements when requesting a route and only nodes that meet these requirements (trust level, metrics), participate in the routing. Questions not addressed by this protocol yet include the mechanism for key distribution, knowledge of the keys of the other nodes, what happens when a node leaves the group with the shared trust level and how trust hierarchies are defined in the first place, especially in civilian applications. SAR relies on tamper-proof hardware.

3.3 Detection, Reputation, and Response Systems

A method for thwarting attacks is prevention. According to Schneier [27], a prevention-only strategy only works if the prevention mechanisms are perfect; otherwise, someone will find out how to get around them. Most of the attacks and vulnerabilities have been the result of bypassing prevention mechanisms. Given this reality, detection and response are essential.

Combining misbehavior detection with a reputation system and appropriate response leads to what we call here a self-policing mobile ad-hoc network. Self-policing means that there are no authorities higher than the nodes themselves. Each node can make their own decisions on how to react to the behavior of other nodes. As opposed to the Byzantine Generals problem, the nodes in a self-policing system for mobile ad-hoc networks do not have to reach a consensus on which nodes misbehave. Each node can keep its own rating of the network denoted by the reputation system entries and it can choose to consider the ratings of other nodes

or to rely solely on its own observations. One node can have varying reputation records with other nodes across the network, and the subjective view of each node determines its actions. Byzantine robustness [22] in the sense of being able to tolerate a number of erratically behaving servers or in this case nodes is the goal of a self-policing system in mobile ad-hoc networks. Here, the detection of malicious nodes by means of observation has to be followed by a response in order to render these nodes harmless.

Since mobile ad-hoc networks have properties that differ from wired networks, such as the lack of infrastructure, misbehavior detection has to be adapted. Every node is their own authority. Nodes can cooperate to compare their notes, but contrary to a wired organized network, one cannot assume that the nodes are under the control of the same organization.

Reputation systems are used to keep track of the quality of behavior of others. In mobile ad-hoc networks, we are interested in the routing and forwarding behavior of nodes. In order to keep track of behavior and to classify it according to whether it is regular or misbehavior for instance, nodes have to be able to observe other nodes. The main goal of reputation systems in mobile ad-hoc networks is to differentiate between regular and misbehaved nodes in order to react accordingly, e.g. by isolating misbehaved nodes from the network.

Only good behavior should pay off in terms of service and reasonable power consumption. Detection of misbehavior has to trigger a response, i.e., a reaction of other nodes that results in a disadvantage for the misbehaved node.

The terms *reputation* and *trust* are being used for various concepts in the literature, also synonymously. We define the term *reputation* here to mean the performance of a principal in participating in the base protocol as seen by others. For

mobile ad-hoc networking this means participation in the routing protocol and forwarding.By the term *trust* we denote the performance of a principal in the policing protocol that aims at protecting the base protocol. For reputation systems this means the reliability as a witness to provide honest reports, in a game-theoretic sense it entails the willingness for retribution, in payment systems the participation in the payment itself.

Self-policing provides a disincentive for cheating by excluding nodes from the network. This isolation also protects the regular nodes. Misbehaved nodes are shunned in two ways. First, nodes route around suspected misbehaved nodes and thus select more reliable routes which increases their throughput. Second, nodes do not provide service to suspected misbehaved nodes, hence their misbehavior ceases to have an impact. The first prevents the misbehaved nodes from being used, the second prevents them from using other nodes.

Reputation systems are not restricted to any one type of misbehaved node, such as selfish, malicious, or faulty.

We now briefly describe some of the protocols proposed in the literature.

Watchdog and path rater components to mitigate routing misbehavior have been proposed by Marti, Giuli, Lai and Baker [15]. They observed increased throughput in mobile ad-hoc networks by complementing DSR with a *watchdog* for detection of denied packet forwarding and a *path rater* for trust management and routing policy rating every path used, which enable nodes to avoid malicious nodes in their routes as a reaction. Ratings are kept about every node in the network and the rating of actively used nodes is updated periodically. Their approach does not punish malicious nodes that do not cooperate, but rather relieves them of the burden of forwarding for others, whereas their messages are forwarded with-

out complaint. This way, the malicious nodes are rewarded and reinforced in their behavior.

CONFIDANT (see our papers [2], [3], [4]) stands for 'Cooperation Of Nodes, Fairness In Dynamic Ad-hoc NeTworks' and it detects malicious nodes by means of observation or reports about several types of attacks and thus allows nodes to route around misbehaved nodes and to isolate them from the network. Nodes have a *monitor* for observations, *reputation records* for first-hand and trusted second-hand observations, *trust records* to control trust given to received warnings, and a *path manager* for nodes to adapt their behavior according to reputation. Simulations for "no forwarding" have shown that CONFIDANT can cope well even with half of the network population misbehaving.

CORE, a collaborative reputation mechanism proposed by Michiardi and Molva [16], also has a *watchdog* component; however it is complemented by a reputation mechanism that differentiates between subjective reputation (observations), indirect reputation (positive reports by others), and functional reputation (task-specific behavior), which are weighted for a combined reputation value that is used to make decisions about cooperation or gradual isolation of a node. Reputation values are obtained by regarding nodes as requesters and providers, and comparing the expected result to the actually obtained result of a request. Nodes only exchange positive reputation information.

A context-aware inference mechanism has been proposed by Paul and Westhoff [20], where accusations are related to the context of a unique route discovery process and a stipulated time period. A combination is used that consists of unkeyed hash verification of routing messages and the detection of misbehavior by comparing a cached routing packet to overheard packets. The decision of how to

treat nodes in the future is based on accusations of others, whereby a number of accusations pointing to a single attack, the approximate knowledge of the topology, and context-aware inference are claimed to enable a node to rate an accused node without doubt. An accusation has to come from several nodes, otherwise a single node making the accusation is itself accused of misbehavior.

OCEAN [1] by Bansal and Baker relies exclusively on first-hand observations. Directly observed positive behavior increases the rating, directly observed negative behavior decreases it by an amount larger than that used for positive increments. If the rating is below the faulty threshold, the node is added to the faulty list. This faulty list is appended to the route request by each node broadcasting it to be used as an avoid list. A route is rated good or bad depending on whether the next hop is on the faulty list. As a response to misbehavior, nodes reject all traffic coming from a suspected misleading node, even if it is not the source of the traffic. The second chance mechanism for redemption employs a timeout after an idle period. Then a node is removed from the faulty list, its rating remaining unchanged. In addition to the rating, nodes keep track of the forwarding balance with their neighbors by maintaining a chip count for each node, which increases when requesting a node to forward a packet and decreases with an incoming request from that node.

3.4 Discussion

Payment systems serve as an incentive to provide a well-defined service, such as packet forwarding, to others for remuneration. The payment has to be unforgeable. To ensure this, tamper-proof hardware and trusted third parties have been suggested. With payment systems, the issue of pricing and other economic questions,

such as how to deal with lost payment, arise. They can prevent selfish forwarding misbehavior, however, they do not address malicious or faulty misbehavior.

Secure protocols prevent preconceived deviations from specific protocol functions. They do, however, not aim at serving as incentives for cooperation or dealing with novel types of misbehavior that occur by going around the protected functions.

Reputation systems apply to a broader range of desired behavior as long as it is observable and classifiable. They can, if they use second-hand information and have means to cope with false accusations or false praise, partially prevent misbehavior by excluding misbehaved nodes. This way, nodes can protect themselves before encountering the misbehaved node. If the reputation systems rely exclusively on first-hand experience to build reputation ratings, they can only prevent more of the misbehavior experienced by a node after it occured.

Preventive schemes can only protect what they set out to protect from the start. There can, however, be unanticipated attacks that circumvent the prevention. It is vital that this misbehavior be detected and prevented from happening again in the future. Self-policing schemes are only as limited as their intrusion detection component regarding detected attacks. The schemes themselves are flexible and can accommodate an evolving intrusion detection component. If the detection of a new attack is conceived of, the detection component can be changed to reflect this added knowledge. This does not in any way change the protocol. If a preventive scheme needs to be extended to accommodate the advent of a new attack, a new version of the routing protocol is required.

As opposed to payment systems, reputation systems do not assume that nodes have to forward for others at least as many packets as they generate themselves.

A self-policing system in the sense of an intrusion detection component with a reputation system merely penalizes a node if it does not do what it is supposed to do according to its own promises. This difference offers an advantage in situations where a node is simply not in the position to cooperate, e.g. when it is at the edge of the network and does not get many requests. In any of the payment systems described here, the node would run out of means to afford having its own packets forwarded by others. This problem is prevented in a self-policing system.

Economic systems assume a rational node that aims at maximizing its utility expressed in power or payment units. The node misbehavior targeted by payment systems is thus selfish concerning utility but it is not malicious.

A malicious node is not necessarily aiming at a economizing on its resources. Its interest lies in mounting attacks on others. Secure routing protocols aim at preventing malicious nodes from mounting attacks.

Although some reactive systems focus on selfish (Watchdog) or malicious misbehavior (intrusion detection), this is not an intrinsic limitation. Self-policing networks can cope with both selfish and malicious, and, in addition, with non intentional fauly misbehavior, the only requirement being that such misbehavior be detectable, i.e. observable and classifiable.

We deem the consideration of non intentional misbehavior such as bugs of high importance, and we think it is vital to protect the network against misbehaved nodes regardless the nature of their intentions. Non intentional misbehavior can result from a node being unable to perform correctly due to a lack of resources, due to its particular location in the network, or simply because of the node being faulty. Self-policing misbehavior detection, reputation, and response systems can be applied irrespective of the actual cause of the misbehavior, be it intentional

or not. When a node is classified as misbehaved it simply means that the node performs badly at routing or forwarding. No moral judgment is implied.

The question of a tamper-proof security module remains controversial [23], but might prove inevitable. As opposed to nuglets and counters, the self-policing reputation systems do not need tamper-proof hardware for themselves, since a malicious node neither knows the entries of its reputation in other nodes nor does it have access to all other nodes for potential modification. The secure module might still be necessary for complementary protection such as authentication.

4 Self-Policing for Mobile Ad-hoc Networks

In this section we explore the properties and trade-offs of self-policing misbehavior detection, reputation, and response systems in more detail. We first describe our own approach to use it as a basis for comparison. We then discuss several issues and contrast the way the proposed approaches address them.

4.1 Enhanced CONFIDANT – a Robust Reputation System Approach

The main properties of a reputation system are the representation of reputation, how the reputation is built and updated, and for the latter, how the ratings of others are considered and integrated. The reputation of a given node is the collection of ratings maintained by others about this node. In our approach, a node i maintains two ratings about every other node j that is cares about. The *reputation rating* represents the opinion formed by node i about node j's behavior as an actor in the base system (for example, whether node j correctly participates in the routing pro-

tocol). The *trust rating* represents node i 's opinion about how honest node j is as an actor in the reputation system (i.e. whether the reported first hand information summaries published by node j are likely to be true).

We represent the ratings that node i has about node j as data structures $R_{i,j}$ for reputation and $T_{i,j}$ for trust. In addition, node i maintains a summary record of *first hand information* about node j in a data structure called $F_{i,j}$.

To take advantage of disseminated reputation information, i.e., to learn from observations made by others before having to learn by own experience, we need a means of incorporating the reputation ratings into the views of others. We do this as follows. First, whenever node i makes a first hand observation of node j's behavior, the first hand information $F_{i,j}$ and the reputation rating $R_{i,j}$ are updated. Second, from time to time, nodes publish their first-hand information to their neighbors. Say that node i receives from k some first hand information $F_{k,j}$ about node j. If k is classified as trustworthy by i, or if $F_{k,j}$ is close to $R_{i,j}$ then $F_{k,j}$ is accepted by i and is used to slightly modify the rating $R_{i,j}$. Else, the reputation rating is not updated. In all cases, the trust rating $T_{i,k}$ is updated; if $F_{k,j}$ is close to $R_{i,j}$, the trust rating $T_{i,k}$ slightly improves, else it slightly worsens. The updates are based on a modified Bayesian approach and a linear model merging heuristic.

Note that, with our method, only first hand information $F_{i,j}$ is published; the reputation and trust ratings $R_{i,j}$ and $T_{i,j}$ are never disseminated.

The ratings are used to make decisions about other nodes, which is the ultimate goal of the entire self-policing system. For example, in a mobile ad-hoc network, decisions are about whether to forward for another node, which path to choose, whether to avoid another node and delete it from the path cache, and whether to

warn others about another node. In our framework, this is done as follows. Every node uses its rating to periodically classify other nodes, according to two criteria: (1) regular/misbehaved (2) trustworthy/not trustworthy. Both classifications are performed using a Bayesian approach, based on reputation ratings for the former, trust ratings for the latter.

Since we apply our reputation system approach to the CONFIDANT [4] protocol, we briefly describe its main features here. The approach we use in CONFIDANT is to find the selfish and/or misbehaved nodes and to isolate them, so that misbehavior will not pay off but result in isolation and thus cannot continue. CONFIDANT detects misbehaved nodes by means of observation or reports about several types of attacks, thus allowing nodes to route around misbehaved nodes and to isolate them.

Nodes have a *monitor* for observations, *reputation records* for first-hand and trusted second-hand observations about routing and forwarding behavior of other nodes, *trust records* to control trust given to received warnings, and a *path manager* to adapt their behavior according to reputation and to take action against misbehaved nodes.

The dynamic behavior of CONFIDANT is as follows. Nodes monitor their neighbors and change the reputation accordingly. If they have reason to believe that a node misbehaves, i.e. when the reputation rating is bad, they take action in terms of their own routing and forwarding. They thus route around suspected misbehaved nodes. Depending on the rating and the availability of paths to the destination, the routes containing the misbehaved node are either re-ranked or deleted from the path cache. Future requests by the badly rated node are ignored. Simulations for “no forwarding” have shown that CONFIDANT can cope well,

even if half of the network population misbehaves.

Note that simply not forwarding is just one of the possible types of misbehavior in mobile ad-hoc networks. Several others, mostly concerned with routing rather that forwarding have been suggested, such as black hole routing, gray hole routing, worm hole routing. Other kinds of misbehavior aim at draining energy, such as the sleep deprivation attack. CONFIDANT is not restricted to handling any particular kind of misbehavior but can handle any attack that is observable. Even if the observation cannot precisely be attributed to an attack but is the result of another circumstance in the network such as a collision, CONFIDANT can make use of it. If it is a rare accident, it will anyhow not influence the reputation rating significantly, and if it happens more often, it means the observed node has difficulties performing its tasks.

4.2 Issues in Reputation Systems for Mobile Ad-hoc Networks

The self-policing systems proposed in the literature differ in several aspects which we explain in the following.

4.2.1 Spurious Ratings

If second-hand information is used to influence reputation, nodes could lie and give spurious rating information. The benefits of false accusations for an adversary are that they can lead to a denial of service of another node by being excluded, false praise can benefit a colluding node. False accusations are not an issue in positive reputation systems, since no negative information is kept [14, 8], however, the disseminated information could still be false praise and result in a good reputation for misbehaved nodes. Moreover, even if the disseminated information is

correct, one cannot distinguish between a misbehaved node and a new node that just joined the network. Many reputation systems build on positive reputation only [26], some couple privileges to accumulated good reputation, e.g. for exchange of gaming items or auctioning [25]. Positive reputation systems are thus used for where one has a choice of transaction partners and wishes to find the best one. In mobile ad-hoc networks, the requirements are different, the focus is on the isolation of misbehaved nodes.

When allowing second-hand information, the question arises whether liars should be punished just as misbehaved nodes are isolated. If we punish nodes for their seemingly inaccurate testimonials, we might end up punishing the messenger and thus discourage honest reporting of observed misbehavior. Note that we evaluate testimonial accuracy according to affinity to the belief of the requesting node along with the overall belief of the network as gathered over time. The accuracy is not measured as compared to the actual true behavior of a node, since the latter is unknown and can not be proved beyond doubt. Even if it were possible to test a node and obtain a truthful verdict on its nature, a contradicting previous testimonial could still be accurate. Thus, instead of punishing deviating views we restrict our system to merely reduce their impact on public opinion. Some node is bound to be the first witness of a node misbehaving, thus starting to deviate from public opinion. Punishing this discovery would be counterproductive, as the goal is precisely to learn about misbehaved nodes even before having had to make a bad experience in direct encounter. Therefore, in our design, we do not punish a node when it is classified as not trustworthy.

4.2.2 Information Dissemination

There is a trade-off between the speed of detection of misbehaved nodes by use of second-hand information and the classification vulnerability introduced by it. CONFIDANT makes use of second-hand information in order to proactively isolate misbehaved nodes before actual encounter. This would make it vulnerable to spurious ratings, notably false accusations, also referred to as blackmailing, and false praise in the case that trusted nodes lie. To prevent that but still retain the advantage of earlier detection, only compatible second-hand information is used and then only slightly influences the reputation rating.

CORE [16] permits only positive second-hand information, which makes it vulnerable to spurious positive ratings and misbehaved nodes increasing each other's reputation. OCEAN [1] relies exclusively on first-hand information for its ratings, trading off detection speed for robustness against spurious ratings. However, it disseminates information about suspected misbhaved nodes by adding them to the avoid list in the route request. Context-aware detection [20] accepts negative second-hand information on the condition that at least four separate sources make such a claim, otherwise the node spreading the information is considered misbehaved. While this distributes the trust given into accusations over several nodes and thus spreads the risk, it inadvertently serves as a disincentive to share ratings and warn others by accusation. Depending on the network density it is also not guaranteed to have at least four witnesses of any event present, let alone four that report it.

4.2.3 Type of Information

The original CONFIDANT used only negative information for the consideration for the reputation system. In the enhanced version as described in [5], also positive information is used to discriminate between active nodes that misbehave sometimes and rarely active nodes that misbehave most of the time. We are thus interested in the relative rate of misbehavior, not the absolute number of misbehavior incidents. OCEAN also uses both positive and negative information for ratings and chip counts. Pathrater and the context-aware detection only consider negative information.

4.2.4 Response

Except for Watchdog and Pathrater [15], all other schemes here have a punishment component in their way of isolating nodes, thus the isolation is twofold: misbehaved nodes are avoided in routes and are denied cooperation when they request it. Not using misbehaved nodes but allowing to be still used by them only increases the incentive for misbehavior, since it results in power saving due to the decrease in number of packets they have to forward for others.

4.2.5 Redemption, Weighting of Time

CORE gives more weight to the past behavior of a node and less to its current behavior. The rationale behind this is that wrong observations or rare behavior changes should not have too much influence on the reputation rating. This holds true only under the assumption that the behavior of a node is constant over time. CONFIDANT takes the opposite approach of discounting the past behavior. This is to ensure that a node can not leverage on its past good performance with its

misbehavior gone unpunished. It also ensures that the system is able to react more quickly to changes of behavior. The other reputation systems do not weight ratings according to time.

Ratings are not only weighted to shift emphasis to the past or the present, but also to add importance to certain kinds of observation. CONFIDANT gives the most weight to first-hand observations and less to reported second-hand information. CORE also uses weights to distinguish between types of observations.

Pathrater, context-aware detection, and OCEAN do not weight ratings according to time.

Redemption has the purpose of mitigating misclassification of a node as misbehaved, either by deceptive observation, spurious ratings, or a fault in the reputation system. Another case that requires redemption is when a node that has been correctly isolated as misbehaved should be allowed back into the network because the root of its misbehavior has been removed, e.g. a faulty node has been repaired, a compromised node has been recaptured by its rightful user.

CONFIDANT allows for redemption of misbehaved or indeed misclassified nodes by reputation fading, i.e. discounting the past behavior even in the absence of in the absence of testimonials and observations, and periodic reevaluation, i.e. checking from time to time whether the rating of a node is above or below the acceptable threshold. Hence, even if a node has been isolated by all nodes, it can get back into the network eventually. Whether it then remains in the network depends on its behavior. Since the ratings do not get erased but only discounted, the rating of the formerly misbehaved node is still close to the threshold value and thus the reaction to renewed misbehavior is swift, resulting in earlier isolation than the misbehavior of a new node. It is thus possible for a node to redeem itself,

given that nodes have each their own reputation belief which is not necessarily shared by all the others. Since their opinions can differ, a node is most probably not excluded by all other nodes and can thus partially participate in the network with the potential of showing its good behavior. Even if this is not the case and the suspect is excluded by everyone it can redeem itself by means of the reputation fading.

In CORE an isolated node should get redemption if it behaves well again, but since it cannot prove itself when isolated, it remains isolated unless there is a sufficient number of new nodes arriving in the network that have no past experience with the isolated node.

OCEAN, like the initial version of CONFIDANT, relies on a timeout of reputation. The sudden lapse back into the network can pose a problem if several nodes set the timer at roughly the same time.

Pathrater and context-aware detection have no notion of redemption.

4.2.6 Weighting of Second-Hand Information

The schemes that use second-hand information have to administer trust of the witnesses, i.e. the sources of second-hand information, in order to prevent blackmailing attacks. The initial CONFIDANT weighted second-hand information according to the trustworthiness of the source and by setting a threshold that had to be exceeded before taking second-hand information into consideration. Second-hand information had to come from more than one trusted source or several partially trusted sources, or any combination thereof provided that trust times number exceeds the trust threshold. This adds a vulnerability of trusting untrustworthy nodes. The notion of trust has been more specifically defined in the enhanced ver-

sion of CONFIDANT, where it means a consistent good performance as a witness, measured as the compatibility between first and second-hand information. This dynamic assessment allows to keep track of trustworthiness and to react accordingly. If the second-hand information is accepted it still only has a small influence on the reputation rating. More weight is given to own direct observation.

The other schemes have no trust management component.

4.2.7 Detection

Reputation systems require a tangible object of observation that can be categorized as good or bad. In online auction or trading systems this is the sale transaction with established and measurable criteria such as delivery or payment delay. For reputation systems on misbehavior in mobile ad hoc networks the analogy to a transaction is not straightforward due to the limited observability and detectability in a mobile and, even more importantly, wireless environment. In order to detect misbehavior, which translates into being able to classify the behavior node as regular, i.e. according to the protocol, or misbehaving, i.e. deviating from the protocol, nodes promiscuously overhear the communications of their neighbors. The component used for this kind of observation is called Watchdog [15], Monitor [4], or Neighbor Watch [1].

The function most used to implement the detection component in the proposed reputation systems is passive acknowledgment [13], where nodes register whether their next-hop neighbor on a given route has attempted to forward a packet. Assuming bidirectional links, a node can listen to the transmissions of a node that is within its own radio range. If within a given time window a node hears a retransmission of a packet by the next-hop neighbor it has sent the packet to previously,

the behavior is judged to be good. Note that this does not necessarily mean that the packet has been transmitted successfully, since the observing node cannot know what goes on outside of its radio range, e.g. there could still be a collision on the far side of the next-hop neighbor.

Several problems with watchdogs have been identified in [15], such as the difficulty of unambiguously detecting that a node does not forward packets in the presence of collisions or in the case of limited transmission power.

In addition to a watchdog-like observation, in CORE nodes do not only rely on promiscuous mode, but in addition they can judge the outcome of a request by rating end-to-end connections [16].

CONFIDANT uses passive acknowledgment not only to verify whether a node forwards packets, but also as a means to detect if a packet, e.g. a routing control message, has been illegitimately modified before forwarding.

4.2.8 Identity.

The question of identity is central to reputation systems. They ideally can assume three properties of identity which we call persistent, unique, and distinct. The requirement to be persistent means that a node cannot easily change its identity. One way of achieving this is by expensive pseudonyms, another is to have a security module. Identity persistence is desirable for reputation systems to enable them to gather the behavior history of a node. An identity is unique if no other node can use it and thus impersonate another node. One way to ensure this is the use of cryptographically generated unique identifiers, as proposed by Montenegro and Castelluccia [18]. This property is needed to ensure that behavior observed was indeed that of the node observed. The requirement of distinct identities is the

target of the so-called Sybil attack analyzed by Douceur [9], where nodes generate several identities for themselves to be used at the same time. This property does not so much concern the reputation system itself, since those identities that exhibit misbehavior will be excluded, while other identities stemming from the same node will remain in the network as long as they behave well. The Sybil attack can, however, influence public opinion by having its rating considered more than once. In the scenario where the mobile ad-hoc network is not completely cut off the Internet, we can make use of certification authorities. An example for such a scenario are publicly accessible wireless LANs with Internet connection. The detection and isolation of misbehaved nodes as achieved by a distributed reputation system for mobile ad-hoc networks are still necessary, even in the presence of network operators.

5 Conclusions

Mobile ad-hoc routing and forwarding are vulnerable to misbehavior, which can occur due to selfish, malicious, or faulty nodes. Solutions to the problem of misbehavior have so far been classifiable into three main categories: payment systems, secure routing, and detection and reputation systems. Payment systems target selfish misbehavior by providing economic incentives for cooperation. Secure routing proposals aim at the prevention of malicious misbehavior. Self-policing systems that consist of detection, reputation, and response components target at the isolation of misbehaved nodes regardless of the reason for misbehavior. None of these solution approaches alone can do prevention, detection, and response. A combination, however, for example of self-policing systems with secure routing can be beneficial to obtain a prevention mechanism along with the advantage of detecting

selfish and faulty misbehavior and providing an adaequate response.

References

[1] Sorav Bansal and Mary Baker. Observation-based cooperation enforcement in ad hoc networks. Technical Report, 2003.

[2] Sonja Buchegger and Jean-Yves Le Boudec. IBM Research Report: The Selfish Node: Increasing Routing Security in Mobile Ad Hoc Networks. RR 3354, 2001.

[3] Sonja Buchegger and Jean-Yves Le Boudec. Nodes Bearing Grudges: Towards Routing Security, Fairness, and Robustness in Mobile Ad Hoc Networks. In *Proceedings of the Tenth Euromicro Workshop on Parallel, Distributed and Network-based Processing*, pages 403 – 410, Canary Islands, Spain, January 2002. IEEE Computer Society.

[4] Sonja Buchegger and Jean-Yves Le Boudec. Performance Analysis of the CONFIDANT Protocol: Cooperation Of Nodes — Fairness In Dynamic Ad-hoc NeTworks. In *Proceedings of IEEE/ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHOC)*, Lausanne, CH, June 2002. IEEE.

[5] Sonja Buchegger and Jean-Yves Le Boudec. A robust reputation system for mobile ad-hoc networks. EPFL Technical Report No. IC/2003/50, July 2003.

[6] Levente Buttyán and Jean-Pierre Hubaux. Enforcing service availability in mobile ad-hoc wans. In *Proceedings of IEEE/ACM Workshop on Mobile Ad Hoc Networking and Computing (MobiHOC)*, Boston, MA, USA, August 2000.

[7] Levente Buttyán and Jean-Pierre Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. Technical Report DSC/2001/046, EPFL-DI-ICA, August 2001.

[8] Chrysanthos Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 150–157, 2000.

[9] John R. Douceur. The sybil attack. In Proc. of the IPTPS02 Workshop, Cambridge, MA (USA), March 2002.

[10] Yih-Chun Hu, David B. Johnson, and Adrian Perrig. SEAD: secure efficient distance vector routing for mobile wireless adoc networks. In *Proceedings of the 4th IEEE Workshop on Mobile Computing Systems & Applications (WMCSA 2002), IEEE, Calicoon, NY, to appear.*, June 2002.

[11] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Ariadne: A secure on-demand routing protocol for adoc networks. Technical Report Technical Report TR01-383, Department of Computer Science, Rice University, December 2001.

[12] Dave B. Johnson and David A. Maltz. The dynamic source routing protocol for mobile ad hoc networks. Internet Draft, Mobile Ad Hoc Network (MANET) Working Group, IETF, Version 9, April 2003.

[13] John Jubin and Janet D. Tornow. The darpa packet radio network protocols. In *Proceedings of the IEEE, 75(1)*, pages 21–32, January 1987.

[14] Peter Kollock. The production of trust in online markets. *Advances in Group Processes, edited by E. J. Lawler, M. Macy, S. Thyne, and H. A. Walker*, 16, 1999.

[15] Sergio Marti, T.J. Giuli, Kevin Lai, and Mary Baker. Mitigating routing misbehavior in mobile ad hoc networks. In *Proceedings of MOBICOM 2000*, pages 255–265, 2000.

[16] Pietro Michiardi and Refik Molva. CORE: A collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. Sixth IFIP conference on security communications, and multimedia (CMS 2002), Portoroz, Slovenia., 2002.

[17] Pietro Michiardi and Refik Molva. Simulation-based analysis of security exposures in mobile ad hoc networks. European Wireless Conference, 2002.

[18] G. Montenegro and C. Castelluccia. Statistically unique and cryptographically verifiable(sucv) identifiers and addresses. NDSS'02, February 2002., 2002.

[19] Peng Ning and Kun Sun. How to misuse aodv: A case study of insider attacks against mobile ad-hoc routing protocols. 4th Annual IEEE Information Assurance Workshop, West Point, June 2003.

[20] Krishna Paul and Dirk Westhoff. Context aware inferencing to rate a selfish node in dsr based ad-hoc networks. In *Proceedings of the IEEE Globecom Conference*, Taipeh, Taiwan, 2002. IEEE.

[21] Charles E. Perkins, Elizabeth M. Royer, and Santanu Das. Ad hoc on demand distance vector (AODV) routing. Rfc 3561, IETF, July 2003.

[22] Radia Perlman. Network layer protocols with byzantine robustness. PhD. Thesis Massachussetts Institute of Technology, 1988.

[23] Andreas Pfitzmann, Birgit Pfitzmann, Matthias Schunter, and Michael Waidner. Trusting mobile user devices and security modules. In *Computer*, pages 61–68. IEEE, February 1997.

[24] P.Papadimitratos and Z.J. Haas. Secure routing for mobile ad hoc networks. In *Proceedings of SCS Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS 2002), San Antonio, TX*. IEEE, January 27-31, 2002.

[25] Paul Resnick and Richard Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. Working Paper for the NBER workshop on empirical studies of electronic commerce, 2001.

[26] Paul Resnick, Richard Zeckhauser, Eric Friedman, and Ko Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.

[27] Bruce Schneier. *Secrets and Lies. Digital Security in a Networked World*. John Wiley $ Sons, Inc, 1 edition, 2000.

[28] William Stallings. *Network and Internetwork Security*. IEEE Press, 2 edition, 1995.

[29] Seung Yi, Prasad Naldurg, and Robin Kravets. Security-aware ad-hoc routing for wireless networks. MobiHOC Poster Session, 2001.

[30] Yongguang Zhang and Wenke Lee. Intrusion detection in wireless ad-hoc networks. In *Proceedings of MOBICOM 2000*, pages 275–283, 2000.

[31] S. Zhong, Y. Yang, and J. Chen. Sprite: A simple, cheat-proof, credit-based system for mobile ad hoc networks. Proceedings of Infocom, 2003.