The iTrust Local Reputation System for Mobile Ad-Hoc Networks

Wei Dai[1]**, L. E. Moser**[1]**, P. M. Melliar-Smith**[1]**, I. Michel Lombera**[1] **and Y. T. Chuang**[1]
[1]Department of Electrical and Computer Engineering, University of California, Santa Barbara, Santa Barbara, CA, USA

Abstract— *The iTrust search and retrieval network serves as a trustworthy medium for the distribution of information that addresses the problems of censorship and filtering of information. To combat subversive behavior of the nodes that might undermine the trustworthiness of iTrust, a reputation system is needed. The iTrust reputation system presented in this paper detects and blacklists malicious nodes. It minimizes the expectation of cooperation between nodes through local reputations based solely on direct observations of the nodes. Simulation results demonstrate that local neighborhoods provide better malicious node detection and blacklisting than does the entire network, which is particularly appropriate for mobile ad-hoc networks.*

Keywords: search and retrieval; mobile ad-hoc network; peer-to-peer network; reputation management; iTrust

1. Introduction

Mobile ad-hoc networks (MANETs) are intrinsically dependent on cooperation and collaboration. MANETs do not rely on a static network infrastructure, but they do rely on several assumptions [18]. Due to the lack of infrastructure and other limiting factors, such as transmission range, symbiotic relationships develop between the nodes in a MANET. Such relationships assume that all of the participants are equally trustworthy and have the same objectives. Such assumptions about the participants are not appropriate for the iTrust search and retrieval system [1], [14], [15], which aims to ensure freedom from censorship and filtering of information, even in the presence of malicious nodes.

A MANET requires cooperation among the nodes in the network to function properly. Without the fulfillment of this requirement, packets would not be forwarded, routes would not be established, and the network would not function properly. Despite the importance of cooperation among the nodes in a MANET, it is not guaranteed. Consequently, a reputation system is needed. However, the addition of a reputation system, in which reports of misbehavior are collected and redistributed, treads dangerously close to encroaching on the fundamental principle of iTrust, which is to provide a distributed, uncensorable, reliable, trustworthy system with no central authority.

The iTrust reputation system, presented in this paper, is based on local reputations and neighborhoods, and uses direct observations of the nodes to detect malicious neighbors, with as few interactions between the nodes as possible. It avoids reliance on information from peers, while still maintaining a method of detecting the misbehavior of malicious or selfish peers. The iTrust reputation system is designed specifically for iTrust MANETs.

In designing the reputation system for iTrust MANETs, we investigated the merits of utilizing a local neighborhood for each node. Simulation results provide increased insight into the rationale behind using local neighborhoods for iTrust. They reveal a distinct relationship between neighborhood size and the number of transmissions required to detect malicious behavior. Essentially, with smaller numbers of transmissions, local neighborhoods consistently yield a higher proportion of malicious nodes detected and blacklisted, compared to the entire network with more transmissions. This finding is particularly important for MANETs, as it is important to eliminate malicious nodes as quickly as possible with as few interactions and transmissions as possible, thereby reducing the costs associated with a reputation system.

The rest of this paper is organized as follows. Section 2 presents an overview of the iTrust search and retrieval network. Section 3 provides an overview of the iTrust reputation system, and the details of its three modules. Section 4 presents an evaluation of the iTrust reputation system, and insight into the use of neighborhoods. Section 5 discusses other reputation systems, and their relationship to the iTrust reputation system. Section 6 concludes the paper and presents future work.

2. The iTrust Search and Retrieval Network

The iTrust search and retrieval network [1], [14], [15] addresses potential problems with centralized search and retrieval systems that are subject to censorship, filtering, and suppression of information. Moreover, the iTrust search and retrieval network is intended to be robust against malicious attacks. To achieve these objectives, iTrust adopts a probabilistic, distributed, and decentralized approach.

The nodes that participate in an iTrust network are referred to as the *participating nodes* (Figure 1). Some of the participating nodes, the *source nodes*, produce information,

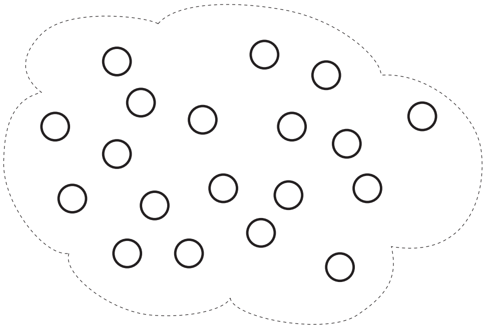

Fig. 1: An iTrust network with participating nodes.

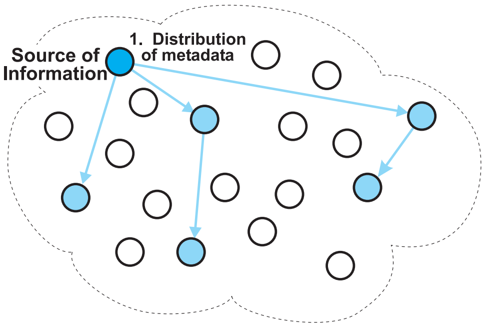

Fig. 2: A source node distributes metadata, describing its information, to randomly chosen nodes in the network.

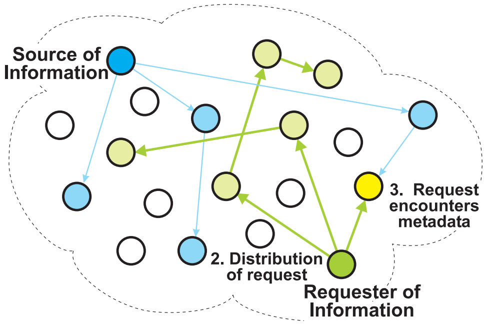

Fig. 3: A requesting node distributes its request to randomly chosen nodes in the network. One of the nodes has both the metadata and the request and, thus, an encounter occurs.

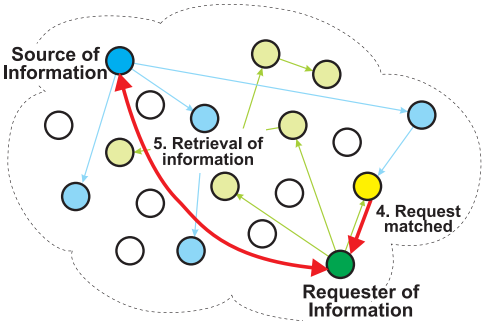

Fig. 4: A node matches the metadata and the request and reports the match to the requesting node. The requesting node then retrieves the information from the source node.

and make that information available to other participating nodes (Figure 2). The source nodes also produce metadata that describes their information, and distribute the metadata, along with the address of the information, to randomly chosen nodes in the iTrust network. Other participating nodes, the *requesting nodes*, request and retrieve information. The requesting nodes generate requests (queries) that contain keywords, and distribute their requests to randomly chosen nodes in the iTrust network (Figure 3). Nodes that receive a request compare the keywords in the request with the metadata they hold. If a node finds a match, which we call an *encounter*, the matching node returns the address of the associated information to the requesting node (Figure 4). The requesting node then uses the address to retrieve the information from the source node. A *match* between the keywords in a request received by a node and the metadata held by a node can be an exact match or a partial match, or can correspond to synonyms.

The iTrust search and retrieval system is based on the hypergeometric distribution [8], given in terms of the following variables:

- n: The number of participating nodes
- x: The proportion of the n participating nodes that are operational, *i.e.*, 1 - x is the proportion of non-operational or malicious nodes
- m: The number of participating nodes to which the metadata are distributed
- r: The number of participating nodes to which the requests are distributed
- k: The number of participating nodes that report matches to a requesting node.

In iTrust, the probability $P(k \geq 1)$ that a request yields one or more matches is given by:

$$P(k \geq 1) = 1 - \frac{n - mx}{n}\frac{n-1-mx}{n-1}\dots\frac{n-r+1-mx}{n-r+1} \quad (1)$$

for $n \geq mx+r$. If $mx+r > n$, then $P(k \geq 1) = 1$. In [14], we showed that, if $m = r = 2\lceil\sqrt{n}\rceil$, then the probability that a request yields one or more matches is $P(k \geq 1) \geq 1-e^{-4} \sim 0.9817$. We use that result and Equation (1) in our evaluation of the iTrust reputation system given in Section 4.

3. The iTrust Local Reputation System

The iTrust local reputation system for MANETs monitors packet forwarding, and watches for non-operational nodes and nodes that do not respond to requests (queries). A local reputation system reduces overheads and the dependence among nodes. It also reduces the amount of storage required, because only information about one-hop neighboring nodes needs to be recorded. In contrast, a global reputation system would result in higher overheads, and also a higher expectation of cooperation among nodes [1].

The iTrust reputation system is based on a *local neighborhood* of each node, consisting of nodes within one hop of the node, and a neighborhood watch mechanism that monitors the interactions of the neighboring nodes. The iTrust reputation system maintains *reputation ratings* of the nodes. A node uses only direct observations to update the reputation ratings of its neighboring nodes. Consequently, the reputation ratings of different nodes might not be consistent. This design choice limits the expectation of cooperation among nodes, thus reducing the opportunities for malicious behavior.

The two primary types of bad behavior that the iTrust reputation system addresses are:

- Malicious behavior: A node does not send responses to requests
- Selfish behavior: A node sends requests and responds to requests, but does not forward messages.

Thus, the iTrust reputation system primarily serves to ensure that peers send messages as expected; it does not address other threats such as Sybil attacks.

In the extreme case in which a node becomes isolated due to the lack of any well-behaved neighbors, the node needs to move to another location where well-behaved nodes are present.

The two main principles under which the iTrust reputation system operates are:

- Intermittent behavior is not punished as much or as rapidly as consistently bad behavior, because intermittent bad behavior is more difficult to detect.
- Efforts are directed towards observing the behavior of nodes within one hop. Malicious behavior that occurs beyond that range is the responsibility of other nodes.

Each node in the MANET maintains a *local reputation table* that consists of a list of nodes within its local neighborhood. Whenever an interaction with another node occurs, the node increases or decreases the reputation rating of that other node. This mechanism addresses malicious or selfish behavior.

The iTrust reputation system consists of three modules that interact with each other. These three modules are the Neighborhood Module, the Reputation Rating Module, and the Monitor Module, which are illustrated in Figure 5 and are described below.

3.1 Neighborhood Module

The Neighborhood Module at a node maintains the local neighborhood of the node and the reputation table for the neighborhood. All of the nodes within one hop of the node, together with their reputation ratings, are represented in the reputation table. Each time a new node is within one hop of the node, the Neighborhood Module adds an entry for the node to the reputation table. A new node starts with a neutral reputation rating of zero. The reputation rating of a node depends on positive and negative interactions with neighboring nodes, as determined by the Reputation Rating Module.

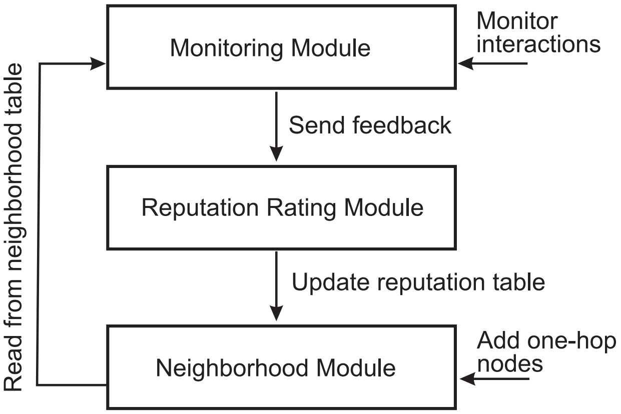

Fig. 5: The three modules of the iTrust reputation system and their interactions.

3.2 Reputation Rating Module

The Reputation Rating Module at a node performs the calculations required to update the node's reputation table. It relies heavily on the Monitoring Module to supply feedback, so that it can decide whether to increase or decrease a node's reputation rating. The Reputation Rating Module is responsible for blacklisting and graylisting nodes.

Blacklisting involves recording malicious nodes in the reputation table. Whenever the reputation rating of a node falls below a certain threshold, the node is blacklisted. A blacklisted node is effectively permanently removed from the neighborhood and from the reputation table maintained by the Neighborhood Module. As a precautionary measure, there is the option of graylisting a node.

Graylisting is a second-chance mechanism that provides a modicum of leniency in an otherwise unforgiving system. Essentially, a node on the graylist is given a second chance before it is blacklisted. A node on the graylist functions as a normal node with the ability to send and receive messages and, to a certain extent, redeems itself through "good" behavior. If a node is graylisted twice, it is put on the blacklist.

3.3 Monitoring Module

The Monitoring Module at a node provides first-hand observations of the behaviors of the nodes within the node's neighborhood. The Monitoring Module provides feedback to the Reputation Rating Module about the good and bad behaviors of neighboring nodes.

For an iTrust MANET, nodes are expected to distribute messages to nodes in the network. Whenever a node interacts with another node by sending a request or a response, it listens to the node's transmissions to check that it is sending messages appropriately. If a node appears to be unresponsive or forwards messages improperly, the Monitoring Module provides feedback to the Reputation Rating Module regarding the negative behavior of the node. The Reputation Rating Module then decreases the reputation rating of the node accordingly.

If a node exhibits malicious or selfish behavior, the Reputation Rating Module decreases the reputation rating of the offending node by -2. If the node exhibits good behavior, the Reputation Rating Module increases the reputation rating of the node by +1. Decreasing a maliciously behaving node's reputation by -2 allows the system to reward good behavior by +1, thus preventing a node from entering cycles of good and bad behavior to dupe a neutral reputation. While we selected -2 and +1 as integers near zero, other values of the reputation rating may be used, as long as the ratio of punishments to rewards is 2:1. The difference in the values of ratings attributed to bad and good behavior allows the system to implement the second-chance mechanism presented in Section 3.2. The reputation rating of a node never exceeds the neutral reputation rating of 0, which is the value it is given when it joins the network. We limit the reputation rating at 0, to prevent a malicious node from building a positive reputation rating over time and then committing a series of malicious acts.

Depending on the level of strictness desired, the user can place a threshold of -2 or -4 on blacklisting a node. With the threshold of -2, a node's first offense results in its being blacklisted. With the threshold of -4, a node is allowed two offenses before it is blacklisted; in this case, a node's first offense results in its being graylisted. These thresholds provide for immediate blacklisting and a second-chance mechanism using graylisting before the node is blacklisted.

With blacklisting for one offense (-2 threshold), a maliciously behaving node is allowed no leniency. If the node behaves maliciously, the reputation system at a node immediately adds the node to the blacklist and permanently bans the node from the node's local neighborhood. In this case, the node's reputation rating is 0 or -2. If the reputation rating is 0, the node sends metadata, requests and responses as usual. If the node's reputation rating falls to -2, the node is blacklisted and effectively removed from the network permanently.

With blacklisting for two offenses (-4 threshold), the reputation system does not blacklist the node immediately; rather, it decreases the reputation rating of the offending node, and places the node on the graylist. If the node behaves maliciously a second time, the system then places the node on the blacklist, and permanently bans the node from the node's local neighborhood. In this case, the node's reputation rating is 0, -1, -2, -3, or -4. If the node's reputation rating is 0 or -1, the node sends metadata, requests and responses as usual. If the node's reputation rating falls to -2, the node is graylisted (or blacklisted if it was previously graylisted). If the node's reputation rating falls to -3 or -4, the node is blacklisted.

In our evaluation of the iTrust reputation system, we investigate the performance of the system with blacklisting for one offense and also blacklisting for two offenses.

4. Evaluation

To evaluate the iTrust reputation system, we perform a simulation. In the simulation, we assume that good behavior can be distinguished from malicious behavior. The experimental setup comprises a network of 1000 nodes, where each node has a neighborhood of 150 nodes. We investigate the advantages and disadvantages of a 150 node neighborhood vs. a 1000 node network, in maintaining the reputations of the nodes and in detecting and blacklisting malicious nodes.

As a baseline, first we investigate the behavior of iTrust in finding a match for the 150 node neighborhood and the 1000 node network. We let m be the number of nodes to which the metadata are distributed, and r be the number of nodes to which the requests are distributed, in the 150 node neighborhood. Similarly, we let M be the number of nodes to which the metadata are distributed, and R be the number of nodes to which the requests are distributed, in the 1000 node network. We investigate the match probability $P(k \geq 1)$ for several values of m and r in the 150 node neighborhood and several values of M and R in the 1000 node network.

4.1 Results without Malicious Nodes

In the first experiment, we set the number M of nodes to which a node in the 1000 node network distributes the metadata to $M = 64$ nodes. Thus, on average, only $m = 9 \sim (64/1000) \times 150$ nodes in the 150 node neighborhood receive the metadata. We set the number R of nodes to which a node in the 1000 node network distributes its requests to $R = 64$ nodes. Likewise, we set the number of nodes to which a node in the 150 node neighborhood distributes its requests to $r = 64$ nodes in the 150 node neighborhood. As shown in Figure 6, the match probability $P(k \geq 1)$ is slightly higher for the 150 node neighborhood than for the 1000 node network.

We then perform two more experiments in which we change the value of r in the 150 node neighborhood. In both experiments, we retain $m = 9$ in the neighborhood, because $m = 9$ represents the proportion of nodes in the 150 node neighborhood that receive the metadata distributed by a node in the 1000 node network.

Thus, in the second experiment, we set $r = 9$ for the 150 node neighborhood, where the value of m is still $m = 9$. As shown in Figure 7, the match probabilities for the 150 node neighborhood are significantly worse than the match probabilities for the 1000 node network with $M = 64$, $R = 64$. We conclude that, to utilize a smaller local neighborhood, it is necessary to increase the number r of nodes in the local neighborhood to which a request is distributed.

Consequently, in the third experiment, we select a value of r for the 150 node neighborhood between 9 and 64. In particular, we choose $r = 24 \sim 2\sqrt{n}$, where $n = 150$, the number of nodes in the local neighborhood. The value of m is still $m = 9$. In Figure 8, it can be seen that the 1000 node network slightly outperforms the smaller 150 node neighborhood, but $r = 24$ is an improvement over $r = 9$ for the 150 node neighborhood. However, there still might be reason to choose the smaller local neighborhood for detecting and blacklisting malicious nodes, which we investigate below.

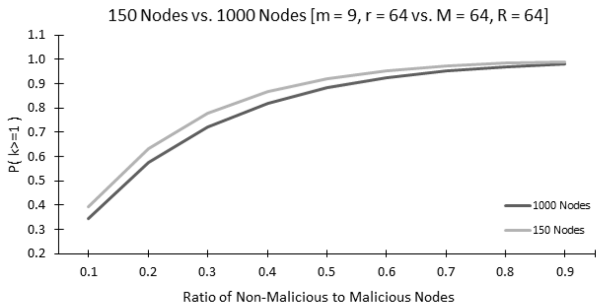

Fig. 6: Probability of one or more matches for the 150 node neighborhood with $m = 9$, $r = 64$ vs. probability of one or more matches for the 1000 node network with $M = 64$, $R = 64$.

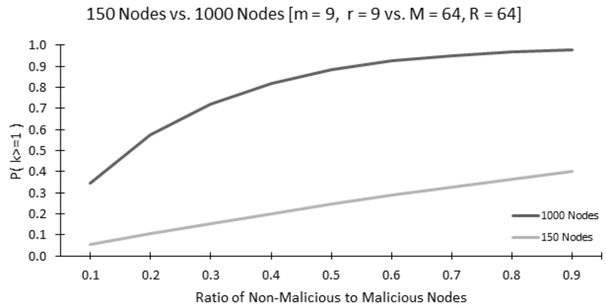

Fig. 7: Probability of one or more matches for the 150 node local neighborhood with $m = 9$, $r = 9$ vs. probability of one or more matches for the 1000 node network with $M = 64$, $R = 64$.

4.2 Results with Malicious Nodes

We now consider malicious nodes within the iTrust network, by having the simulator randomly flag nodes as malicious. For these experiments, we use a proportion of 0.2 malicious nodes, or 200 malicious nodes in the 1000 node network. Thus, the 150 node neighborhood contains, on average, $30 = (200/1000) \times 150$ malicious nodes.

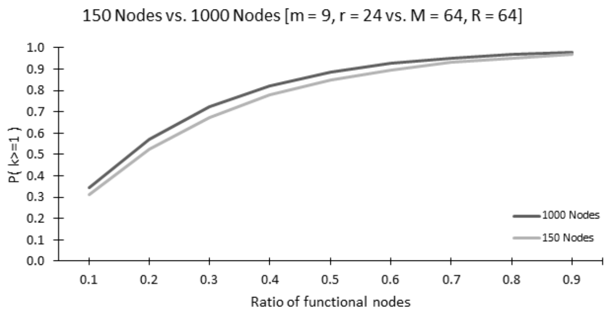

Fig. 8: Probability of one or more matches for the 150 node neighborhood with $m = 9$, $r = 24$ vs. probability of one or more matches for the 1000 node network with $M = 64$, $R = 64$.

If a requesting node's message is mishandled by a malicious node, the requesting node detects the malicious node, and either blacklists the offending node immediately, or allows it to have a second chance by placing it on the graylist. We investigate both possibilities, and run the simulation for 10, 100, 1000, and 10000 requests.

The results are presented in the following tables and graphs. For example, when Requests equals 10, there are 10 sets of metadata, each of which is distributed to m nodes in the local neighborhood and to M nodes in the entire network, and there are 10 requests, each of which is distributed to r nodes in the local neighborhood and to R nodes in the entire network. The Proportion Blacklisted is calculated as Blacklisted/30 for the 150 node neighborhood and as Blacklisted/200 for the 1000 node network.

In the first experiment, we investigate the use of blacklisting for two offenses in the 150 node neighborhood with $m = 9$, $r = 64$, and in the 1000 node network with $M = 64$, $R = 64$. As shown in Table I, the proportion of malicious nodes blacklisted varies with the number of requests and the number of nodes. With 10000 requests, both the 150 node neighborhood and the 1000 node network are successful in blacklisting a similar proportion of malicious nodes, 0.90 and 0.91, respectively. However, when the number of requests is 1000, we see significant differences in the proportion of malicious nodes blacklisted. The 150 node neighborhood is still successful in blacklisting 0.87 of the malicious nodes, but the 1000 node network is able to detect merely 0.13 of the malicious nodes. For 100 requests, the difference is even greater. The 150 node neighborhood recognizes 0.70 of the malicious nodes, but the 1000 node network detects none. In terms of the number of requests it takes for malicious nodes to be blacklisted for two offenses, a smaller local neighborhood performs better than a larger network, as is evident in Figure 9.

In the next experiment, we investigate the use of blacklisting for one offense. The values of the parameters are the same as those in the previous experiment. In Table II,

Nodes	Distribution	Requests	Blacklisted	Remaining	Proportion Blacklisted
150	$m = 9$	10	3	27	0.10
	$r = 64$	100	21	9	0.70
		1000	26	4	**0.87**
		10000	27	3	**0.90**
1000	$M = 64$	10	0	200	0.00
	$R = 64$	100	0	200	0.00
		1000	25	175	**0.13**
		10000	182	18	**0.91**

Table 1: The 150 node neighborhood with $m = 9$, $r = 64$ vs, the 1000 node network with $M = 64$, $R = 64$, with blacklisting for two offenses.

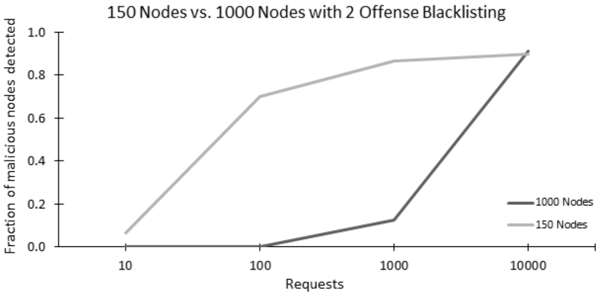

Fig. 9: Proportion of malicious nodes blacklisted for two offenses for various numbers of requests for the 150 node neighborhood with $m = 9$, $r = 64$ and for the 1000 node network with $M = 64$, $R = 64$, with blacklisting for two offenses.

we see that with blacklisting for one offense, the 150 node neighborhood still outperforms the 1000 node network for 10000 requests. At 1000 requests, a drop in the proportion blacklisted occurs, similar to the previous case with blacklisting for two offenses. While less severe, the difference between the 150 node neighborhood and the 1000 node network is still significant, with 0.93 vs. 0.34 of the malicious nodes blacklisted. Overall, we see an increase in the proportion of malicious nodes blacklisted when compared with blacklisting for two offenses. However, the proportion blacklisted might be inflated by the detection of false positives. Blacklisting for one offense is more severe, and does not take other factors into account, such as temporary loss of connectivity. Doing so leads to more non-malicious nodes being blacklisted, making blacklisting for two offenses a more reasonable choice, despite the marginal improvement in the proportion blacklisted, compared to blacklisting for one offense. As Figure 10 shows, the difference between the 150 node neighborhood and the 1000 node network is still significant, even with blacklisting for one offense.

The final experiment that we performed investigates the 150 node neighborhood with $r = 24 \sim 2\sqrt{150}$ requests distributed, compared to the 1000 node network, with blacklisting for two offenses and blacklisting for one offense, as shown in Table III and Table IV, respectively. Whereas

Nodes	Distribution	Requests	Blacklisted	Remaining	Proportion Blacklisted
150	$m = 9$	10	6	24	0.20
	$r = 64$	100	24	6	0.80
		1000	28	2	**0.93**
		10000	26	4	0.87
1000	$M = 64$	10	3	197	0.02
	$R = 64$	100	10	190	0.05
		1000	68	132	**0.34**
		10000	183	17	0.92

Table 2: The 150 node neighborhood with $m = 9$, $r = 64$ vs. the 1000 node network with $M = 64$, $R = 64$, with blacklisting for one offense.

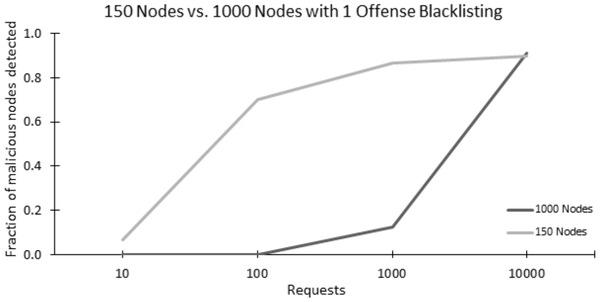

Fig. 10: Proportion of malicious nodes blacklisted for the 150 node neighborhood with $m = 9$, $r = 64$ vs. the 1000 node network with $M = 64$, $R = 64$, with blacklisting for one offense.

a 150 node neighborhood leads to slightly lower match probabilities than does a 1000 node network, in terms of finding malicious nodes, the 150 node neighborhood is able to catch more malicious nodes, for a given number of requests, with both blacklisting for two offenses and blacklisting for one offense.

Nodes	Distribution	Requests	Blacklisted	Remaining	Proportion Blacklisted
150	$m = 9$	10	0	30	0.00
	$r = 24$	100	8	22	0.27
		1000	25	5	0.83
		10000	29	1	0.97
1000	$M = 64$	10	0	200	0.00
	$R = 64$	100	0	200	0.00
		1000	25	175	0.13
		10000	182	18	0.91

Table 3: The 150 node neighborhood with $m = 9$, $r = 24 \sim 2\sqrt{150}$ vs. the 1000 node network with $M = 64$, $R = 64$, with blacklisting for two offenses.

Table V aggregates the results presented previously, so that we can compare the performance of blacklisting for two offenses and blacklisting for one offense, for the 150 node neighborhood and the 1000 node network. Again we consider $r = 24 \sim 2\sqrt{150}$, $m = 9$ for the 150 node

Nodes	Distribution	Requests	Blacklisted	Remaining	Proportion Blacklisted
150	$m = 9$	10	3	27	0.10
	$r = 24$	100	21	9	0.70
		1000	27	3	0.90
		10000	29	1	0.97
1000	$M = 64$	10	3	197	0.02
	$R = 64$	100	10	190	0.05
		1000	68	132	0.34
		10000	183	17	0.92

Table 4: The 150 node neighborhood with $m = 9$, $r = 24 \sim 2\sqrt{150}$ vs. the 1000 node network with $M = 64$, $R = 64$, with blacklisting for one offense.

neighborhood, and $M = 64$, $R = 64$ for the 1000 node network. As the table shows, for 10000 requests, the proportions of malicious nodes blacklisted are all greater than 0.9. Moreover, for 1000 requests, the proportions of malicious nodes blacklisted for the 1000 network are substantially less than the proportions of malicious nodes blacklisted for the 150 node neighborhood, with blacklisting for both two offenses and one offense.

	150 Nodes $m = 9$ $r = 24$ 2 offenses	1000 Nodes $M = 64$ $R = 64$ 2 offenses	150 Nodes $m = 9$ $r = 24$ 1 offense	1000 Nodes $M = 64$ $R = 64$ 1 offense
Requests	Proportion Blacklisted			
10	0.00	0.00	0.10	0.02
100	0.27	0.00	0.70	0.05
1000	**0.83**	**0.13**	**0.90**	**0.34**
10000	0.97	0.91	0.97	0.92

Table 5: Proportion of malicious nodes blacklisted as a function of the number of requests for the 150 node neighborhood with $m = 9$, $r = 24 \sim 2\sqrt{150}$ vs. the 1000 node network with $M = 64$, $R = 64$, with blacklisting for two offenses and blacklisting for one offense.

One could also investigate blacklisting for three offenses, although we would expect an even greater decrease in the proportion of malicious nodes blacklisted, particularly for the 1000 node network.

These simulations demonstrate the effectiveness of smaller local neighborhoods in an iTrust MANET. Whereas a traditional reputation system requires repeated interactions between nodes to build a reputation table, iTrust seeks to reduce the number of interactions with its smaller local neighborhoods.

5. Related Work

Cho *et al.* [6] present a survey of trust management for MANETs. They discuss classifications, potential attacks, performance metrics, and in particular a trust metric that combines the notion of trust from social networks with quality-of-service. In [5], Cho *et al.* investigate selfish behavior in packet forwarding within MANETs. Their analysis balances altruism, *i.e.*, forwarding packets for the public good, against selfish individual welfare, *i.e.*., not forwarding packets to conserve battery power; however, it does not consider malicious behavior. Such an analysis might be interesting for the iTrust reputation system, but might be vulnerable to malice.

Damiani *et al.* [7] enumerate a range of malicious behaviors that can distort the reporting of nodes' behaviors and the evaluation of nodes' reputation ratings in the Gnutella peer-to-peer network [9]. Their approach to collecting reputation information is based on gathering reports from large numbers of nodes, and on gathering reports for both the resources and the nodes that provide access to those resources. Their global approach does not fit in with iTrust's local neighborhoods that aim to limit the expectation of cooperation among nodes.

Buchegger and Le Boudec [4] investigate a Bayesian approach to evaluate a node's reputation from second-hand reports obtained from other nodes, which is contrary to iTrust's aim to limit the interactions between nodes. To protect a node's reputation against malicious reports, reputation reports that are inconsistent with the node's current reputation are rejected, which might result in the failure to adjust a node's reputation in the presence of subtle malicious attacks. Extending this work, Mundinger and Le Boudec [17] employ an interesting mean-field approach. Such an approach is effective at masking uncorrelated noise, but might not be able to handle correlated misinformation in coordinated malicious attacks.

Guo *et al.* [10] take a different approach to monitoring packet forwarding in wireless ad-hoc networks. They exploit fuzzy sets with mathematical analysis based on Grey theory to detect inconsistent and potentially malicious behavior. We are investigating whether such an approach can be extended to the rather more complex behavior of iTrust nodes.

Zhou and Hwang [20] present a distributed reputation system that places more weight on nodes considered to be the most reputable. Doing so can result in a system that is dominated by a small number of nodes which, in turn, can result in subtle malicious attacks. To address this issue, Jesi *et al.* [13] aim to detect hub attacks. Because hubs concentrate power over reputations, routing, *etc.* into relatively few nodes, hub attacks can distort or disrupt the behavior of the system.

The Collaborative Reputation (CORE) system, developed by Michiardi *et al.* [16], for MANETs uses a collaborative monitoring technique and reputation mechanism, where reputation is based on a node's ability to cooperate with other nodes. Nodes with good reputations are granted the use of resources, whereas nodes with bad reputations are gradually filtered out. Their watchdog mechanism is similar to the neighborhood watch mechanism of the iTrust reputation system; however, their mechanism is less well protected against malicious manipulation of reputation information.

Jelasity *et al.* [12] maintain logs of all outgoing and incoming messages, with signed messages to preclude forgeries. Periodically, the nodes exchange logs, which allows them to check the behavior of other nodes and to detect various kinds of malicious behavior. Such a strategy is less effective in MANETs, where the neighborhoods change quite quickly as the nodes move around. The iTrust reputation system uses only first-hand observations to monitor the behavior of the neighboring nodes.

Ruohomaa *et al.* [19] developed a peer-to-peer reputation system in which nodes distribute their reputation rankings to other nodes. In their system, potential interactions are described by a collaboration contract. Whether such interactions satisfy the contract is verified using non-repudiation receipts, thus preventing reputations from being distorted by misinformation. Adding such a mechanism to the iTrust reputation system would be quite expensive.

Hu [11] presents a reputation system that resists malicious attacks. Nodes develop reputations of their neighbors from observations of their neighbors' behaviors. Their system does not communicate the reputations to other nodes, thus making it more difficult for a malicious node to subvert the reputations. The iTrust reputation system adopts a similar strategy.

The Cooperation Of Nodes: Fairness In Dynamic Ad-hoc Networks (CONFIDANT) protocol, proposed by Buchegger and Le Boudec [3], attempts to detect and isolate uncooperative nodes. The nodes use passive observation of packets forwarded within a one-hop neighborhood. To prevent dissemination of false reputation ratings, the system incorporates a trust rating for each node. First-hand information is stored locally and disseminated to neighbors, but reputation and trust ratings are not shared. Similarly, the iTrust reputation system does not share reputation information among the nodes.

The Observation-based Cooperation Enforcement in Adhoc Networks (OCEAN) system, developed by S. Bansal and M. Baker [2], recognizes that reporting a node's behavior to other nodes renders the system vulnerable to malicious reports. It focuses on first-hand observations of other nodes' behaviors, exploiting the ability of nodes in wireless adhoc networks to listen to the transmission of neighboring nodes. Simulations demonstrate that OCEAN works quite well, even though ratings are based only on monitoring neighboring nodes. Likewise, the iTrust reputation system adopts a local neighborhood strategy.

6. Conclusions and Future Work

The iTrust local reputation system for MANETs detects malicious nodes and puts such nodes on a blacklist or a graylist. For 10000 requests, the results for a 150 node neighborhood and a 1000 node network are similar, in detecting and blacklisting malicious nodes. In contrast, for fewer requests, the 150 node neighborhood yields superior results to the 1000 node network, with most of the malicious nodes being blacklisted. In a MANET, such as that of iTrust, having a high level of repeated interactions with the same nodes is rare. Therefore, relying on a large number of requests to detect and blacklist malicious nodes is unrealistic; the smaller local neighborhood provides a means to eliminate the need for a large number of interactions between nodes.

The current reputation system for iTrust mitigates the effects of subversive nodes with respect to messages. However, misbehaving nodes are still capable of disseminating bad data. In a future version of iTrust, we plan to incorporate a mechanism that monitors the message content and rates the information at the end user, perhaps with the user's help. As a result, it will also be able to rate the source nodes, based on the content they distribute. This addition will improve the overall robustness of iTrust against malicious behavior.

References

[1] C. M. Badger, L. E. Moser, P. M. Melliar-Smith, I. Michel Lombera and Y. T. Chuang, "Declustering the iTrust search and retrieval network to increase trustworthiness," in *Proc. 8th International Conference on Web Information Systems and Technologies*, Porto, Portugal, Apr. 2012, pp. 312–322.

[2] S. Bansal and M. Baker, "Observation-based cooperation enforcement in ad hoc networks," Tech. Rep. NI/0307012, Stanford Univ., 2003.

[3] S. Buchegger and J. Y. Le Boudec, "Performance analysis of the CONFIDANT protocol (Cooperation Of Nodes: Fairness In Dynamic Ad-hoc NeTworks)," in *Proc. IEEE/ACM Symposium on Mobile Ad Hoc Networking and Computing*, Lausanne, Switzerland, June 2002, pp. 226–236.

[4] S. Buchegger and J. Y. Le Boudec, "A robust reputation system for P2P and mobile ad-hoc networks," in *Proc. Second Workshop on the Economics of Peer-to-Peer Systems*, Harvard University, June 2004.

[5] J. H. Cho, A. Swami and I. R. Chen, "Modeling and analysis of trust management protocols: Altruism versus selfishness in MANETs," in *Proc. IFIP Conference on Trust Management*, June 2010, Morioka, Japan, pp. 141–156.

[6] J. H. Cho, A. Swami and I. R. Chen, "A survey of trust management for mobile ad hoc networks," *IEEE Communications Surveys and Tutorials*, vol. 13, no. 4, 2011, pp. 562–583.

[7] E. Damiani, D. di Vimercati, S. Paraboschi, P. Samarati and F. Violante, "A reputation-based approach for choosing reliable resources in peer-to-peer networks," in *Proc. 9th ACM Conference on Computer and Communications Security*, Nov. 2002, Washington, DC, pp. 207–216.

[8] W. Feller, *An Introduction to Probability Theory and Its Applications I*, John Wiley & Sons, New York, NY, 1968.

[9] Gnutella, 2000, http://gnutella.wego.com/

[10] J. Guo, A. Marshall and B. Zhou, "A trust management framework for detecting malicious and selfish behavior in ad-hoc wireless networks using fuzzy sets and Grey theory," in *Proc. IFIP Conference on Trust Management*, June 2011, Copenhagen, Denmark, pp. 277–289.

[11] J. Hu and M. Burmester, "Cooperation in mobile ad hoc networks," In: *Guide to Wireless Ad Hoc Networks*, Springer, London, 2009, pp. 43–57.

[12] M. Jelasity, A. Montresor and O. Babaoglu, "Detection and removal of malicious peers in gossip-based protocols," in *Proc. 2nd Bertinoro Workshop on Future Directions in Distributed Computing*, University of Bologna, Bertinoro, Italy, June 2004, pp. 23–25.

[13] G. P. Jesi, D. Hales and M. van Steen, "Identifying malicious peers before it's too late: A decentralized secure peer sampling service," in *Proc. First International Conference on Self-Adaptive and Self-Organizing Systems*, Bologna, Italy, July 2007, pp. 237–246.

[14] P. M. Melliar-Smith, L. E. Moser, I. Michel Lombera and Y. T. Chuang, iTrust: Trustworthy information publication, search and retrieval, in *Proc. 13th International Conference on Distributed Computing and Networking,* Hong Kong, China, Jan. 2012, pp. 351–366.

[15] I. Michel Lombera, L. E. Moser, P. M. Melliar-Smith and Y. T. Chuang, “Mobile decentralized search and retrieval using SMS and HTTP,” *ACM Mobile Networks and Applications Journal*, vol. 18, no. 1, 2013, pp. 22–41

[16] P. Michiardi and R. Molva, “Core: A collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks,” in *Proc. IFIP Communication and Multimedia Security Conference*, Canterbury, UK, Sep. 2012, pp. 107–121.

[17] J. Mundinger and J. Y. Le Boudec, “Analysis of a reputation system for mobile ad-hoc networks with liars,” *Performance Evaluation*, vol. 65, no. 3–4, Mar. 2008, pp. 212-âĂŤ226.

[18] E. M. Royer and C. K. Toh, “A review of current routing protocols for ad-hoc mobile wireless networks,” *IEEE Personal Communications Magazine*, Apr. 1999, pp. 46–55.

[19] S. Ruohomaa, P. Kaur and L. Kutvonen, “From subjective reputation to verifiable experiences – Augmenting peer-control mechanisms for open service ecosystems,” in *Proc. IFIP Conference on Trust Management*, May 2012, Surat, India, pp. 142–157.

[20] R. Zhou and K. Hwang, “PowerTrust: A robust and scalable reputation system for trusted peer-to-peer computing, *IEEE Transactions on Parallel and Distributed Systems*, vol. 18, no. 4, Apr. 2007, pp. 460–473.