Trusting Routers and Relays in Ad hoc Networks.

Prashant Dewan and Partha Dasgupta
Arizona State University
dewan@asu.edu, partha@asu.edu

Abstract

The current generation of ad hoc networks relies on other nodes in the network for routing information and for routing the packets. These networks are based on the fundamental assumption that the nodes will cooperate and not cheat. This assumption becomes invalid when the network nodes have tangential or contradictory goals.

A novel method of enhancing routing strategies, and enhancing cooperation is to use "reputations" computed from peer recommendations. Reputation assignment and use cajole cooperation from the nodes in the network even if they do not share the same goal. This paper provides a mechanism that can use reputations in ad hoc network for trusting routers and relays. In addition, it enumerates the issues involved in using reputation in ad hoc networks.

The simulations show that the throughput of the network increases by 0% - 71.6% when the neighbor reputations and shortest path are considered, for deciding the next hop. The throughput of the network improves from 3% to 143% when the next hop of the packet is decided using only reputations and ignoring the shortest path. The average hop length is the same irrespective of the fact that reputations are used.

1. Introduction

Wireless ad hoc networks are self-configuring, adaptive networks, which can be deployed in areas deprived of any existing network infrastructures. Mobile ad hoc networks do not need any infrastructure to deploy these networks rapidly. These networks are generally used in situations of crisis, where the existing network has been mutilated due to natural or organized calamities. Besides, these networks are used in places where the existing infrastructure is not available because of monetary or strategic reasons. For example, they can be used in battle zones, villages and in areas suffering from natural calamities.

The current research on efficient routing in mobile networks [Perkins et al. *1999], [Park et al. *1997] makes the fundamental assumption that all the nodes in the network have the same or, at least, similar goals; and hence it needs to cooperate with other nodes in the network. However this assumption does not hold in certain scenarios. A node of an ad hoc network may be compromised. Compromised nodes become antagonistic to un-compromised nodes; and therefore, they are not reliable for retrieving the routing information or for actual routing.

Another example of ad hoc networks is a network where the entity controlling the nodes has disparate goals. Such a network needs external motivation to make the nodes of the ad hoc network cooperate. In addition, this mechanism should be able to weed out rogue or compromised nodes and inject a certain level of inter–node trust. Network nodes can use this trust to estimate the probability that the packet that they send will eventually reach its destination.

Several incentive mechanisms have been proposed in [Obreiter et al. *2003] where the authors classify the incentive mechanisms as *account based* and *reputation based* mechanisms. Account based mechanisms are based on the assumption that the nodes in an ad hoc network try to accumulate the maximum amount of micro payments and then exchange it for real world currency.

Account based mechanisms do not appear to work in collaborative groups of nodes in the ad hoc network, where the nodes set common goals. The nodes in such collaborative networks try to achieve a common goal and not to make money. Such techniques can be used in non-collaborative ad hoc networks where the network does not have one goal but several with tangential goals belonging to the members of the network. In such situations micro payments can be cashed for real world money and hence motivate the network nodes to cooperate and make money.

The *reputations* of the nodes can be used for instilling the motivation to cooperate in the nodes. This motivation also establishes trust and confidence among the nodes. The reputation of a node also motivates it to act in a trust worthy fashion and not to maliciously tamper with any packet. If a node becomes indifferent to its reputation and continuous to act maliciously, it is weeded out of the network.

Ad hoc networks do not depend on the presence of any centrally trusted authority. Hence there is a need for distributed protocols which facilitates the nodes in an ad hoc network to collect, store and manage the reputations of other nodes. In addition, the nodes should be able to quickly use the reputation information to take routing decisions without having a significant impact on the routing performance.

2. Design Goals

In this research, the following goals were set.

1. The system should be robust to collaboration attacks. If a certain number of nodes are compromised, they should not able to collaborate and bring the whole network down.
2. The reputation information should be easy to use and the nodes should be able to ascertain the best available nodes for routing without requiring human intervention.
3. The system should have a low performance cost because low routing efficiency can drastically affect the efficiency of the applications running on the ad hoc network.
4. Nodes should neither be able to fake nor manipulate their reputations in the system.
5. Nodes should be able to punish the wrong doer by providing him with a bad reputation

3. Overview of Reputations & Routing

A reputation can be defined as an import of the past behavior of an entity. It has been used by the entities, in the brick and mortar world to find out the probability of achieving the desired level of satisfaction from their transaction with the other entity whose reputation is being evaluated. The reputations were used in the networks in general, and Internet in particular, only in the late 1990s. An example of a successful reputation based system is Ebay.

Although in the brick and mortar world, the reputations have been used subjectively, such representation cannot be used in the digital world, like the Ad hoc networks. Hence we assign numerical values to reputations in order to make them objective and hence more usable. A more exhaustive description of reputations and their use in networks in general, and self-configuring networks in particular, can be found in [Dewan et al. *2003].

Routing in ad hoc networks has been well researched. An exhaustive overview of the different routing protocols used in ad hoc networks is provided by [Royer et al. *1999]. The current methods of routing use the "ask the neighbor" method recursively, till they reach the destination or another node, which has the route to the destination. Furthermore these routing techniques are classified into *Table Based Routing*, and *On Demand Routing*. The table based routing is used for updating the routes periodically. The nodes store the routes to all the destinations while the on demand routing is used for updating the routes on demand.

Before sending a packet to its destination, a node has to know, at least, the possible next hops for the destination of the packet. In addition, the node needs to know the metric that shows the cost involved in sending the packet via that path. For example, the path length can be such a metric. This paper proposes the use of reputation of the next hop node besides the path length, to choose the path to be followed by the packet. In other words, the node performs a trade of between the shortest path and the neighbor reputation.

Also, the current routing protocols make the assumption that the nodes will not cheat. ARAN [Sanzgiri et al. *2002] is a security protocol in which makes non-cooperation difficult by using cryptographic techniques like digital certificates.

4. Definitions

In section 3.1 we define some of the more commonly used terms in this paper.

4.1 Recommendation

The recommendation is the value assigned to the service provider by the service seeker during a transaction. All the recommendations of the service provider are combined to evaluate its reputation. Recommendations can be contextual for a given node. For example, a node can get different recommendations for its availability, accuracy, and efficiency.

In this paper, we consider only one context: whether a given node is a "good" router or a "bad" router. In most cases a good node routes the received packets to the next hop even if it has no vested interest in the packet. A bad node maliciously drops the packets or tampers with the contents of the packet or routes it in the wrong direction.

For example, consider a network that consists of three nodes

$$A \rightarrow B \rightarrow C$$

If node A wants to send a message to node C and it finds out that the only way that it can send the message to node C is via B. It sends the message to B, which in turn routes it to C. If C acknowledges receiving the message to A, A can deduce that B routed the message properly and hence gives B a recommendation of +1. Mathematically, the recommendation will appear as follows,

$$Rep_{AB} = +1$$

4.2 Reputation

Mathematically, reputation is the mean of the recommendations received by a node. Suppose node B received 100 packets and routed 80 packets but dropped 20 packets, the senders of routed packets give him +1 and the senders of the dropped packets give him −1. Hence his total reputation becomes

$$R_B = (80-20)/100 = 60/100 = 0.6$$

4.3 Node Identifier

Each node possesses a certificate, which was issued to it when the network was established. No node possesses multiple identities. The reputation is assigned to the only identity that the node possesses. The nodes in the network can easily verify the identity of a particular node in the network by using challenge response

mechanisms.

If a node gets compromised and does not co-operate, its reputation goes down quickly and soon it is ostracized and weeded out of the system, even if it possesses an authentic identity.

4.4 Threshold Reputation

The threshold reputation R_{thresh} is the minimum reputation a node expects from a possible next hop on a path. If the next hop node does not possess the requisite reputation, only then the source node sends a packet to a node whose reputation is less than the threshold.

5. Router Reputations

In the proposed technique all the nodes in the network use the reputations of their neighboring nodes to find out the best node, to which the packet should be forwarded.

5.1 Finding Trusted Routers

Each node in the system possesses a reputation. When it joins the network with a certain identifier the reputation of a node is one, which is the maximum possible reputation. During the course of time, the reputations of the nodes decrease or remain unchanged. They are based on the number of packets forwarded by the node.

We consider three scenarios here for deciding the next hop for a packet. In the first scenario, only the shortest path to the destination is used. In the second scenario, the shortest path to the destination and reputation of the neighbors are used. In the third scenario, only the reputations of the neighbors having a path to the destination are used.

In scenario 2, once the routing information is available, a node chooses the next hop, which provides the shortest path to the destination. At this point we add another factor to the routing decision, reputation. The node sorts all the possible paths to a given destination by using the length of the paths to the destination, which is followed by the reputations of the next hop nodes. The sender then selects the first M available next hops having a reputation greater than the threshold reputation R_{th} from the sorted list. It then sends M copies, numbered from 1 to M of the packet on these paths. The node which receives this packet, sorts its available next hops by shortest path to the destination and reputations and sends it to the first node with a reputation greater than R_{thresh}.

The important thing to note here is that although the first node sends M copies, the subsequent nodes only send one copy each. In addition, M should be as small as possible because an increase in M will lead to a multifold increase in the traffic. In our simulations, we use M=1. The packets have identification numbers called PID and hence any intersecting paths merge. The merging nodes send the list of the numbers of a given packet received by the node.

Once the packet reaches the destination, it makes a list of the numbers of copies that it received and sends it to the sender. After the sender receives the packets, it finds out the next hops of the successful paths; and gives the recommendation of +1 to the next hop node(s) in the network and a recommendation of –1 to the next hop nodes of the paths, which failed. The next hop nodes pass on their recommendations and give a recommendation of –1 to their next hop if the sender gives them –1.

In scenario 3, a node generates the list of its neighbors having a path to the destination. The node sorts the neighbor list on the basis of their reputations. It selects the top M nodes from the list and sends the packet to these M nodes. The sender then selects the first M available next hops having a reputation greater than the threshold reputation R_{th} from the sorted list. It then sends M copies, numbered from 1 to M of the packet on these paths. The node, which receives the packet, sends one copy of the packet to the node having the highest reputation in the set of neighbors having a path to the destination.

In other words, the responsibility of making the packet reach its destination relies on the next hop node of the sender. If the packet is dropped anywhere in the path between the source and the destination, the next hop node of the source gets a bad recommendation which can be cascaded to all the nodes in the path. In order to maintain high reputations, the intermediate nodes in the path, treat the packet as their own and forward it to the next possible hop till it reaches the destination.

5.2 The Incentive

Due to the autonomous nature of the nodes in ad hoc networks, they need incentives to maintain their high reputation levels. The nodes that have reputation levels lower than the threshold are chucked out of the network. The packets sent by nodes with higher reputation levels get a priority over the packets sent by nodes with lower reputation levels. Hence, nodes with higher reputations experience a lower latency compared to nodes with lower reputation assuming all other conditions remain constant.

5.3 Recommendations Management

Managing recommendations in the network is a difficult task, which entails collecting, storing and deducing the reputations from the recommendations. There are two models for managing reputation. In the first model, the recommendations are stored on various nodes of the network; and any node can search for any other node's recommendations using the public key of the other node. In the second model the recommendation receiver stores the recommendations that it receives and presents them to other nodes that want to know the reputation of this node.

The advantage of the first technique is that, it is harder to maliciously manipulate the reputations because they are not stored at the node-which is effected by them. The disadvantage of this technique is that, any node

looking for the recommendation of another node, needs to search of the order of O (log N), where N is the number of nodes in the network. This search can be expensive if the size of the network is large.

The advantage of the second approach is that any node looking for the reputation of another node can get the recommendations directly from the owner node, for example, the node that is affected by the recommendations. The disadvantage of the second approach is that it becomes easier to maliciously modify a recommendation. Hence heavy weight cryptographic mechanisms are needed in order to prevent any malicious tampering.

In the simulation we assume the first model. In addition we assume that the nodes do not falsify their reputation.

A more detailed discussion of both the techniques can be found in [Dewan et al. *2003]. In short, the tradeoff is between the amount of resources available to each node and the potential size of the ad hoc network.

5.4 Discussion

5.4.1 Motivation for Autonomous nodes

The nodes in an ad hoc network are more or less autonomous. Hence this technique motivates them to allocate their resources to other nodes in the network. The ad hoc networks are built upon the assumption that the nodes will cooperate.

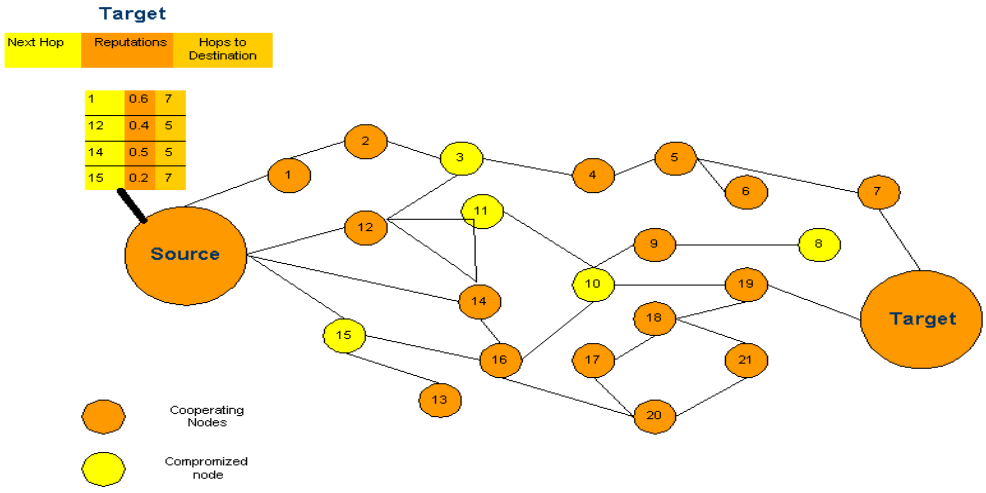

Figure 1

5.4.2 Increased Throughput

As the sender only sends the packet to highly trusted nodes, based on their history, it lowers the risk that its neighbors will intentionally drop the packet. The neighbors in turn forward the packets to nodes trusted by them. As a result the number of intentionally dropped packets is reduced and hence the throughput of the system is higher.

5.4.3 Non-cooperative Nodes are Ostracized

Non-cooperative nodes are slowly weeded out of the network. A node held by an individual who does not want to cooperate with the other nodes in the network or nodes which were earlier cooperating but were compromised in a battle and became non cooperating start accumulating negative reputation.

As a result most of the rogue nodes i.e. the nodes having low reputation do not receive any packet, and hence cannot inflict any damage by dropping packets.

5.4.4 Poor Nodes are penalized

Some nodes in the network might be unable to route packets for other nodes due to the scarcity of resources. Such nodes get lower reputation because of this disability and might be subsequently weeded out for no fault of theirs.

This problem is solved by attaching a list of resources available to the given node in its certificate. This class of nodes is penalized for failing to route packets- but the penalty inflicted on them is a fraction of what is inflicted on a node with a large volume of resources.

For example, consider node A, which possesses half the memory and half the processing power of most of the other nodes in the network. If node A fails to route a packet sent to it by some other node, it gets a recommendation of –0.5 instead of -1. In this way the system is democratized.

5.4.5 Good Nodes become a bottleneck

The good nodes receive more and more packets as their reputation increases. As a result the good nodes are not able to handle the amount of load received by them. Hence they drop some of the packets and their reputation goes down. The senders of the packets use the second rank nodes to route their packets. Slowly the system comes to equilibrium.

5.4.6 PKI is Expensive and Routing is Resource Intensive.

Nodes in a mobile ad hoc network are severely constrained by resources like memory and power. They may not have the resources to perform expensive cryptographic operations like PKI. Even if the protocol is useful some of the nodes might not have the resources to execute it. This is one major bottleneck that needs to be circumvented

6. Algorithm

In this algorithm we assume that the nodes already know their neighbors. The neighbors of a node are defined as the nodes within the transmission range of the source node. Neighbor discovery can be done via broadcasting or by asking other neighbors. Initially the reputations of all the neighbors are 1, i.e. good.

This algorithm is independent of the routing scheme used. The only requirement is that the current node should know the route to the destination from itself. It also assumes that the nodes store their own recommendations and the recommendations are not stored by any central agency. Before executing this protocol, the node carrying the packet collects the reputations of its neighbors. The reputations are only based on the previous routing history of the node and do not consider the possibility of malicious tampering of the packet by uncooperative nodes.

7. Example

This example emulates scenario 2. In figure 1, the Source wants to send a packet to the Target. The value of M, number of paths established by the source with the target is 2 and the minimum amount of reputation expected, R_{thresh}=0.5. The nodes 3,8,10,11,15 are rogue or compromised nodes, which have a history of non-cooperation. In other words their reputation is below the threshold reputation.

Using any routing protocol, it finds out that its possible next hops for reaching the Target is 1, 12,14, and 15. Out of the 4 possible neighbors, 1 and 15 need 7 hops to reach the target while 12 and 14 need 5 hops to reach the target. The node sorts the list on the basis of the number of hops required to reach the target. The sorted table is shown as Table 1.

```
Let Source be S and Destination be D
Let the current node where the packet is be denoted by C

S-> Copy the source node, S to the current node, C.
While (C <> D)
Do
  Step 1 C->Find Route to Destination using any standard routing technique like AODV.

Step 2 If C is uncooperative drop the packet


  Step 3

(Scenario 2) C-> Find the first neighbor in the shortest route to the destination whose reputation is above the threshold, R_thresh

(Scenario 3) C-> Find the neighbor in the set of routes to the destination having the best reputation.

  Step 4 If the neighbor is found go to Step 1

Step 5 If the neighbor is not found, sends it to the uncooperative neighbor of the highest reputation.
Done
```

```
If ACKNOWLEGEMENT Received
  S->Grant +1 to the NEIGHBOR

IF No ACKNOWLEDGEMENT Received
S->Grant -1 to the NEIGHBOR

Update Reputation of the Neighbor
```

Algorithm

Next Hop	Reputation	Hops to Destination
12	0.4	5
14	0.5	5
1	0.6	7
15	0.2	7

Table 1

It selects nodes 14 and 12. It sends the packet to node 14 since it provides the shortest path at a reputation above the threshold. It needs to select one more next hop from nodes 12,1 and 15. It does not select 12 because its reputation is below the threshold level. For the same reason it does not select 15. It selects node 1. It makes two copies of the packet and numbers them. The first packet is sent to node 14 and the second packet is sent to node 1.

Node 14 sends it to node 16, which routes it to 20 and subsequently to the target. If reputation had not been used it would have sent it to 10 and because 10 is a rogue node, the packet would have never reached the target.

The packet, which was sent to 1, reaches 3 via 2. As node 3 is a rogue node, it drops the packet. Once the target receives the packet, it acknowledges the packet number to the source. The source finds out that the packet sent to 1 never reached. Hence it gives a recommendation of –1 to node 1 and +1 recommendation to 14. The recommendations get cascaded and nodes 2 and 3 get a recommendation of –1 while 14,16,17,18,19,20and 21 get a recommendation of +1. The update reputation table is shown in Table 2

Next Hop	Old Reputation	New Reputation	Hops to Destination
12	0.4	0.4	5
14	0.5	0.75	5
1	0.6	-0.2	7
15	0.2	0.2	7

Table 2

8. Simulation

8.1 Setup

The simulation is performed on a 2.4 GHZ, Linux Red Hat machine using C language. Two scenarios are simulated. The first scenario is routing performed without using the reputation scheme that is proposed in this paper, and the second scenario is the routing performed with the reputation scheme. The third scenario is routing performed without calculating the shortest path and using the reputation scheme.

The main focus of the simulation is to compare the throughputs of the system in all the scenarios i.e. compare the number of packets that reached their destination and the number of packets dropped, in the above scenarios. The second important parameter is the average number of hops in the network in the above scenarios.

The simulator is designed as an N*N matrix of nodes where N=100, and L= 500 links between them. The links are randomly generated. Each node keeps track of its neighborhood nodes and its reputation. The value of M, the number of copies of the same packet sent by a node, is set to 1. Hence if a packet reaches a node whose all neighbors for a given route are not cooperative, the packet is sent to one of the uncooperative nodes, which subsequently drop it. The threshold reputation is set to 0.7. Any node whose reputation falls below threshold during the simulation is considered to be uncooperative. P=1000 packets are sent from randomly selected sources to randomly selected destinations.

Initially it is assumed that all the nodes are cooperative. In the subsequent simulations more nodes are made uncooperative till the number of cooperative and uncooperative nodes is equal. If the packet reaches its destination, all the intermediate nodes get a recommendation of +1. If the packet is dropped then all the nodes in the route before the node that dropped the packet get a recommendation of –1. The reputation of a node is the average of the recommendations received by it.

One important negative effect of reputations, which has not been measured in this simulation, is the number of nodes whose reputation went below the threshold, because of the fact that they did not have sufficient resources to forward packets. Table 3 shows the important parameters

8.2 Simulation Analysis

The following salient observations are made from the results of the simulation.

1. The number of packets reached is always more in when the reputations of neighbors is used for routing. When the number of uncooperative nodes is zero then the numbers of packets that reached their destination is same in both scenarios.
2. Using reputations along with shortest path information to the destination increases the throughput of the system to almost 71% when 50% of the nodes are rouge. This is because the packet is now routed via cooperative nodes, which do not have a history of dropping the packets. Hence the number of packets that reach their destination is higher.
3. When reputations are used for the selection of the next hop without using the shortest path information, the number of packets that reach their destination increases by 143% when 50% of the nodes are rouge.
4. The increase in the throughput is explained by the following example. Consider three nodes Alpha, Beta, Gamma. Beta is the shortest path from Alpha to Gamma. Here Beta is initially a good node but turns uncooperative. Supposed Alpha has to send 5 different messages to Gamma. When Alpha sends the first message, Beta drops it.

Hence Beta's reputation reduces below the maximum. But its reputation is still above the threshold. In scenario 2, Alpha keeps on sending the packets to Beta till its reputation goes below the threshold, as it is the shortest route to Gamma. Hence all the packets are dropped till the reputation of Beta goes below the threshold and no more packets are sent to Beta.

In contrast, in scenario 3, after the first packet is dropped and the reputation of Beta goes down maximum, Alpha does not send any more packets to Beta as long as other routes are available. Hence the number of packets that are dropped reduces.

5. The percentage of packets that reach their destination, as compared to the packets sent without the uncooperative nodes is 35%. Although this value is very low for ad hoc networks, this figure can be attributed to the sparseness of the simulated network. Only 500 links are present out of a possible of 49,950 links. It is only 0.1% of the number of links in a fully connected network. Hence it must have resulted in a large number of networks, which are not interconnected. Such a scenario is very much possible in mobile ad hoc networks.
6. Such a sparsely connected network has been used in order to estimate the impact in worse scenarios. A network with higher number of links will obviously display a better throughput. The average number of hops, sum of the number of hops traversed by successful packets divided by the total number of successful packets. It is the same whether reputations are or are not used for routing.
7. The average number of hops is reduced when the number of uncooperative nodes has increased. This can be attributed to the fact that the number of packets reaching their destination is reduced. The packets whose source and destination have less number of nodes are less likely to find uncooperative nodes in their path than the packets, which have longer path lengths.

Number of Nodes, N=100, Number of Links, L=500, Number of Packets, P=1000						
	Scenarios	Rouges	Reached	Dropped	Average Number of Hops	Throughput Increase (%)
1.	1	0	353	647	2	-
2.	2	0	353	647	2	0
3.	3	0	365	635	2	3.3
4.	1	10	285	715	2	-
5.	2	10	300	700	2	5.2
6.	3	10	317	683	2	5.9
7.	1	20	194	806	2	-
8.	2	20	220	780	2	13.4
9.	3	20	254	746	2	30.9
10.	1	30	158	842	2	-
11.	2	30	190	810	2	20.2
12.	3	30	223	777	2	41.1
13.	1	40	108	892	2	-
14.	2	40	139	861	2	28.7
15.	3	40	180	820	2	66.6
16.	1	50	53	937	1	-
17.	2	50	91	909	1	71.6
18.	3	50	129	871	1	143.3

Table 3

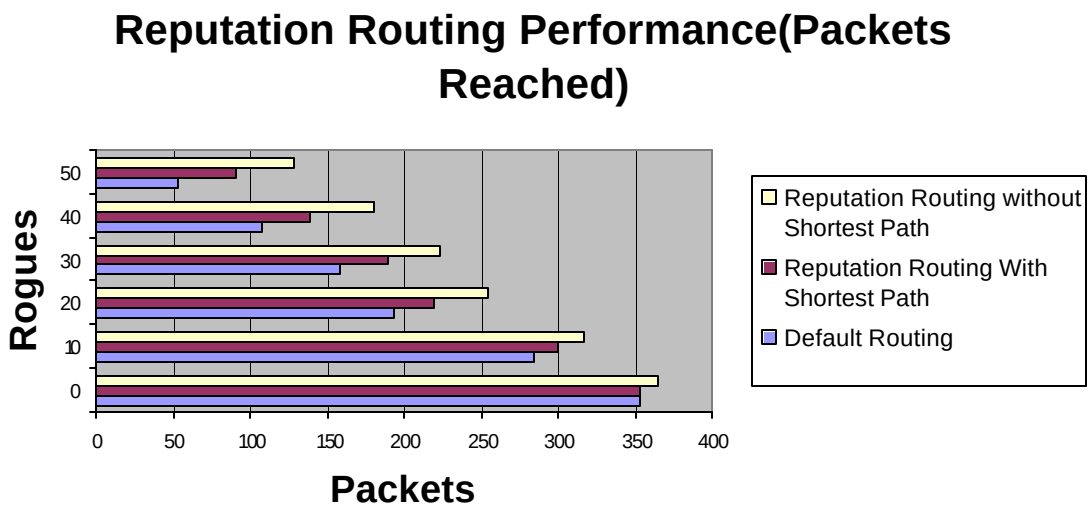

Figure 2

9. Conclusion

Ad hoc networks provide a communication medium where infrastructure networks are absent. Ad hoc networks rely on cooperation of the network nodes for routing. In the absence of a common goal the nodes need an external motivation to cooperate.

Reputations of the nodes and the subsequent advantages associated with having high reputation can provide the motivation for the node to commit their own resources to others. Using reputations with shortest path information increases the throughput from 3 to 72% over routing by using only the shortest path information. Using reputations without shortest path increases the throughput up to 143%.

Reputation systems need techniques for collecting and storing recommendations and making sure that the rogue nodes cannot misrepresent their reputation. Another set of challenges is posed by the resource constraints of the ad hoc network nodes. Encryption and addition work performed for routing can overload the nodes and hence raise overall latency. In other words, this extra load will provide increased throughput and security at the cost of more work and higher latency.

References

[Dewan et al. *2003] Prashant Dewan, Austin Godber, and Partha Dasgupta. 2003. A Self-Certification Scheme for Managing Reputations in Peer-to-Peer Networks. (submitted)

[Obreiter et al. *2003] Philipp Obreiter, Birgitta Koenig-Ries, and Michael Klein. 2003. Stimulating Cooperative Behavior of Autonomous Devices - An Analysis of Requirements and Existing Approaches. Paper read at Second International Workshop on Wireless Information Systems (WIS2003), April 22,2003, at Angers, France.

[Park et al. *1997] Vincent D. Park, and M. Scott Corson. 1997. A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks. Paper read at INFOCOM.

[Perkins et al. *1999] Charles E. Perkins, and Elizabeth M. Royer. 1999. Ad-hoc On-Demand Distance Vector Routing. Paper read at 2nd IEEE Workshop on Mobile Computing Systems and Applications, February 1999, at New Orleans.

[Royer et al. *1999] Elizabeth M. Royer, and C-K Toh. 1999. A Review of Current Routing Protocols for Ad-Hoc Mobile Wireless Networks. *IEEE Personal Communications Magazine*, April 1999, 46-55.

[Sanzgiri et al. *2002] Kimaya Sanzgiri, Bridget Dahill, Brian Neil Levine, Clay Shields, and Elizabeth M. Belding-Royer. 2002. A Secure Routing Protocol for Ad Hoc Networks. Paper read at International Conference for Networking Protocols, November 12-15, 2002.