Reputation Methods for Routing Security for Mobile Ad Hoc Networks*

Po-Wah Yau and Chris J. Mitchell
Mobile VCE Research Group
Information Security Group
Royal Holloway, University of London
Egham, Surrey TW20 0EX, UK

P.Yau@rhul.ac.uk, C.Mitchell@rhul.ac.uk

ABSTRACT

Mobile ad hoc networks have inherently very different properties to conventional networks. These new characteristics present major security vulnerabilities; in particular, one side effect of the unique way in which routing protocols operate in ad hoc networks is that many new threats arise. Selfish nodes are those which do not perform certain operations that the protocol specifies that they should, through a wish to conserve power. Malicious nodes may deliberately disrupt the network using a variety of attacks. This paper discusses reputation mechanisms which have been suggested as a means to mitigate the detrimental effect of selfish and malicious nodes. This paper reveals reasons why complex reputation systems may be too inefficient to use in a mobile ad hoc network, where resources are limited. However, suggestions are also made to show how a simple reputation system might be used to enhance the robustness of ad hoc networks.

Keywords

Routing protocols, network security, mobile networks.

1. INTRODUCTION

Mobile ad hoc networks have inherently very different properties to conventional networks. An ad hoc network is a collection of nodes forming a temporary or permanent network without any support from centralised services. Within a wireless network, a node's transmission range will typically not cover the whole network, so end-to-end communication may require routing information via several nodes. This is why ad hoc networks are sometimes referred to as multi-hop networks, where a hop is a direct link between two nodes. Such systems have a variety of security issues, many of which are different to the issues surrounding conventional wired networks. Reputation systems have been suggested as a tool to resolve some of the security issues associated with ad hoc network routing. This paper discusses the effectiveness of applying reputation systems to mobile ad hoc networks, where resources can be limited.

Section 2 introduces the security issues in ad hoc networks. Section 3 give an overview of reputation systems and how they have been applied to ad hoc networks. Section 4 discusses the use of reputation systems, analysing how reputations are calculated and how they are distributed. Section 5 suggests how simple reputation systems can be used to help improve the security of ad hoc networks. Finally, section 6 concludes the paper.

The following terms are used in this document, but may be used differently elsewhere. A *node* is a device which has a network interface participating in routing in a mobile ad hoc network. It may or may not be mobile, and may also be part of another network. It is important to realise that a node can actually be a large network, or it could just be a single mobile device such as a mobile phone. An *originator node* is a node which originates a data packet, intended for a certain *destination node*. A node is a *neighbour node* of another node if it is only one hop away and within direct transmission range. If the destination node is not a neighbour node of the originator node, the data packet will have to traverse a multi-hop route consisting of *intermediate nodes*. In a specific scenario, the *sending node* is the last node to have forwarded the data packet. A *service* is defined as an action that a node requests from, or performs for, its neighbour node.

2. SECURITY ISSUES

The lack of infrastructure management in an ad hoc network gives problems with the provision of any security services, which are typically centrally controlled in conventional wired networks. For example, access control is a service traditionally maintained by a central server, controlling various resources belonging to the network nodes. In an ad hoc network context, such a service may be required to help prevent unauthorised principals interfering with a private network.

Security mechanisms involving trusted third parties may no longer be viable in ad hoc networks. As nodes are mobile, continually entering and leaving the ad hoc network, a dynamic topology means that security will have to be scalable and cope with frequent link breaks. This will be of particular importance in safety applications, where availability is key. For example, if safety-critical data is to be sent, then it is imperative that the information is reliably and speedily delivered. As communication uses wireless technology, bandwidth will often be limited, as may be transmission en-

*The work reported in this paper has formed part of the Networks & Services area of the Core 2 Research Programme of the Virtual Centre of Excellence in Mobile & Personal Communications, Mobile VCE, www.mobilevce.com, whose funding support, including that of EPSRC, is gratefully acknowledged. Fully detailed technical reports on this research are available to Industrial Members of Mobile VCE

ergy. These constraints introduce issues with heterogeneous networking, where resource-intensive security mechanisms may not work in an ad hoc environment.

One of the key research areas in mobile ad hoc networks is setting up and maintaining the ad hoc infrastructure through the use of routing protocols. Existing network layer protocols are likely to be too resource intensive to be suitable for ad hoc network use, so solutions using a variety of different methods have been proposed. One side effect of the unique way in which these routing protocols work is that many more threats now exist. In an earlier paper [9], a threat model for ad hoc routing protocols was described, classifying internal attacks into four categories — failed, badly failed, selfish, and malicious nodes. Failed and selfish nodes are those which do not perform certain operations that the protocol specifies that they should, the former due to some unforeseen failure and the latter due to selfishness to conserve power. Badly failed nodes may perform operations incorrectly, introducing false and misleading information into the network. Malicious nodes may deliberately disrupt the network using a variety of attacks.

3. REPUTATION MECHANISMS

Reputation mechanisms, the main focus of this paper, have been proposed for use within ad hoc networks to address some of the threats arising from misbehaving network nodes. These mechanisms, explored in more detail immediately below, are potentially of particular value in addressing the threats arising from selfish nodes. In the context of an ad hoc network, these mechanisms seek to dynamically assess the trustworthiness of neighbouring network nodes, with a view to excluding untrustworthy nodes.

The use of reputation systems in many different areas of IT is increasing[1], not least because of their widely publicised use in online auctions and product reviews, see, for example eBay and Amazon [8]. Mui et al. [6] give many examples of how reputation systems are used.

Reputation systems are used to decide who to trust, and to encourage trustworthy behaviour. Resnick and Zeckhauser [7] identify three goals for reputation systems:

1. To provide information to distinguish between a trustworthy principal and an untrustworthy principal,
2. To encourage principals to act in a trustworthy manner, and
3. To discourage untrustworthy principals from participating in the service the reputation mechanism is present to protect.

Reputation systems rely on principals monitoring sequences of transactions with other principals, and on communications between principals that are willing to take part in the reputation system. Each principal maintains a reputation value for some subset of the other principals in the system — these values may be shared between principals or may be unique for each participant. The precise meaning of the reputation value, how it is calculated and updated, and how it is communicated between parties, are all system-dependent. However, it is generally true that this value is intended in some way to measure the trustworthiness of the principal, at least for the purposes of the system concerned.

Two reputation mechanisms that have been proposed to help protect ad hoc routing are the Cooperation of Nodes: Fairness in Dynamic Ad-Hoc NeTworks (CONFIDANT) protocol [1], and the Collaborative Reputation Mechanism (CORE) protocol [5], which work in a similar way. These two schemes are now described in more detail.

3.1 An Overview of CORE

Michiardi and Molva [5] define an ad hoc network as a community where each member has to contribute to its running to remain trusted. Any member not contributing will find their reputation worsening until they are gradually excluded from the operation of the network because of their bad reputation.

3.1.1 Reputation Values

CORE defines three types of reputation, which are combined to form a global reputation value for a community member. Each calculation is normalised so that a reputation ranges from –1 (bad) to +1 (good). 0 represents a neutral view, and this is used when there are not enough observations to make an accurate assessment of a node's reputation. The three reputation types are as follows:

- *Subjective reputation* is locally calculated, where node A calculates the reputation of a neighbour node B at a given time for a particular function. More emphasis is given to past behaviour than current behaviour. Placing more weight on past observations prevents subjective reputation being influenced by sporadically correct behaviour.
- *Indirect reputations* are accepted by node A from node C about node B. Only positive reputation values are used, to eliminate an attack where a malicious node transmits negative reputation information to cause a denial-of-service.
- *Functional reputations* are related to a certain function, where each function is given a weight as to its importance. For example, data packet forwarding may be deemed to be more important than forwarding packets with route information, so data packet forwarding will be given greater weight in the reputation calculations.

Reputation values in CORE are based on observations. If the observed behaviour matches the expected behaviour then the k-th observation will be positive. If not then the k-th value is negative. To be able to perform this validation reliably is of extreme importance to the CORE scheme, and the authors have suggested the Watchdog mechanism [4] based on promiscuous observation[2]. The expected result is stored in a buffer until a matching observation is made. While the expected result is still present in the buffer, the reputation relating to the observed function is gradually decreased.

[1]The Reputations Research Network is a web site for discussing reputation system research at `databases.si.umich.edu/reputations/index.html`

[2]There are inherent problems with promiscuous observation which are discussed in section 4.2.

3.1.2 The CORE Protocol

Each node maintains a reputation table. This table consists of the reputations of other nodes, with each entry consisting of a unique ID, recent subjective observations, recent indirect observations and the composite reputation for a given function. Thus a reputation table has to be maintained for each function that is to be monitored.

There are three ways in which a reputation table is updated.

1. In the first case, a node A requests a service from node B, but node B refuses to perform the service. Thus node A will decrease its perceived reputation of node B. This is a local calculation of node B's subjective reputation.

2. In the second case, a global distribution of reputation takes place within a reputation dissemination phase. This phase involves sending messages containing a list of entities which have successfully cooperated in providing a function, i.e. a list of nodes with positive reputations.

3. The final case is where reputations are gradually decreased to a null value if there is no interaction with the observed node.

When a node A, with a good reputation, is asked to perform a service by a node B, who has a bad reputation, node A can refuse to cooperate. In doing so, node A is required to send a message to all nodes in the ad hoc network, stating that it is denying service to node B. The neighbour nodes of both A and B must check that node B's reputation is negative in their own reputation tables. If one of the neighbour nodes does not agree with node A's negative reputation value for node B, then this neighbour node decreases the reputation of node A, i.e. the node which sent the denial-of-service message.

3.2 An Overview of CONFIDANT

The CONFIDANT protocol [1] is defined as a collection of components at a node which interact with each other to provide and process protocol information. See figure 1 for an overview of how these components interact.

Nodes in the CONFIDANT scheme rely on passive observation of all packets within a one-hop neighbourhood. This view is maintained by the Monitor component. The Monitor reports any suspicious events to the Trust Manager component.

The Trust Manager makes decisions about providing or accepting route information, accepting a node as part of a route, or taking part in a route originated by another node. The Trust Manager maintains a trust table indicating how much other nodes in the network can be trusted to send correct ALARM messages. ALARM messages, which contain the type and frequency of protocol violations, are recorded in an alarm table. ALARM messages are sent to all nodes in a 'friends' list whenever bad behaviour is experienced or observed, or when other valid ALARM messages have been received from trusted nodes. The level of trust assigned to a node in the trust table is used as a weighted metric to determine whether or not an ALARM message is credible, e.g. two ALARMS from partially trusted nodes are equivalent to one ALARM from a trusted node.

The Monitor and Trust Manager report suspicious events, either from an ALARM message or direct observation, to the Reputation System component. This component maintains either a table of nodes and their reputation ratings[3], or a blacklist of nodes which the component believes are behaving badly (or both). These ratings can only be changed if there is sufficient evidence to indicate a change is needed, i.e. evidence has been received at least a threshold number of times. Evidence is weighted, in order from highest to lowest, depending whether it results from personal experience, from observed experience in the neighbourhood, or from reported ALARM messages. If a rating drops out of a certain range, the reputation system alerts the Path Manager.

The Path Manager ranks routes according to a security metric. All paths which contain a badly behaving node are deleted. The Path Manager also decides what to do with route requests received from badly behaved nodes, e.g. by ignoring the bad node, and what to do with route requests for a bad node, e.g. it could alert the originator node.

4. THE APPLICABILITY OF REPUTATION SYSTEMS

The previous two schemes are examples of reputation systems as applied to an ad hoc network environment. However, the reputation system itself potentially introduces many new vulnerabilities. Quantifying reputation is very difficult and must be defined in a very precise way, and distributing reputation information in a reliable and secure manner is particularly difficult in ad hoc networks.

4.1 Analysis of the Reputation Value

Reputation can be defined as one node's perception of another node with regard to performing some operation. Thus the reputation value is used as a prediction of future quality of service. However, reputation is not a tangible property, so many different definitions exist. Thus the reputation value has to be explicitly defined.

4.1.1 Calculating Reputation

Reputation values are inevitably based on observations of multiple functions. In both CORE and CONFIDANT, weightings are applied to these individual functional reputation values to obtain a single 'combined' reputation value. The CORE mechanism assumes that every node will use the same reputation calculations and will also assign the same weights to the same functions. This is a potentially inappropriate assumption in a multi-domain ad hoc network, where devices with different capabilities and roles are likely to want to place different levels of importance on different functions.

Both CORE and CONFIDANT use global reputation values, i.e. each node maintains a single reputation value for every other node with which it interacts, where this value combines all the various functional reputation values. Issues arise with use of such global reputation values. In particular, a global reputation value may enable a node to hide bad behaviour with respect to one function by correctly supporting

[3]It is implicitly implied that this table is independent of the similar table maintained by the Trust Manager. There is no explicit definition of what a rating is and what value/s it takes. The authors state that a node can be good or bad, but they give no detailed definition on what is good and what is bad.

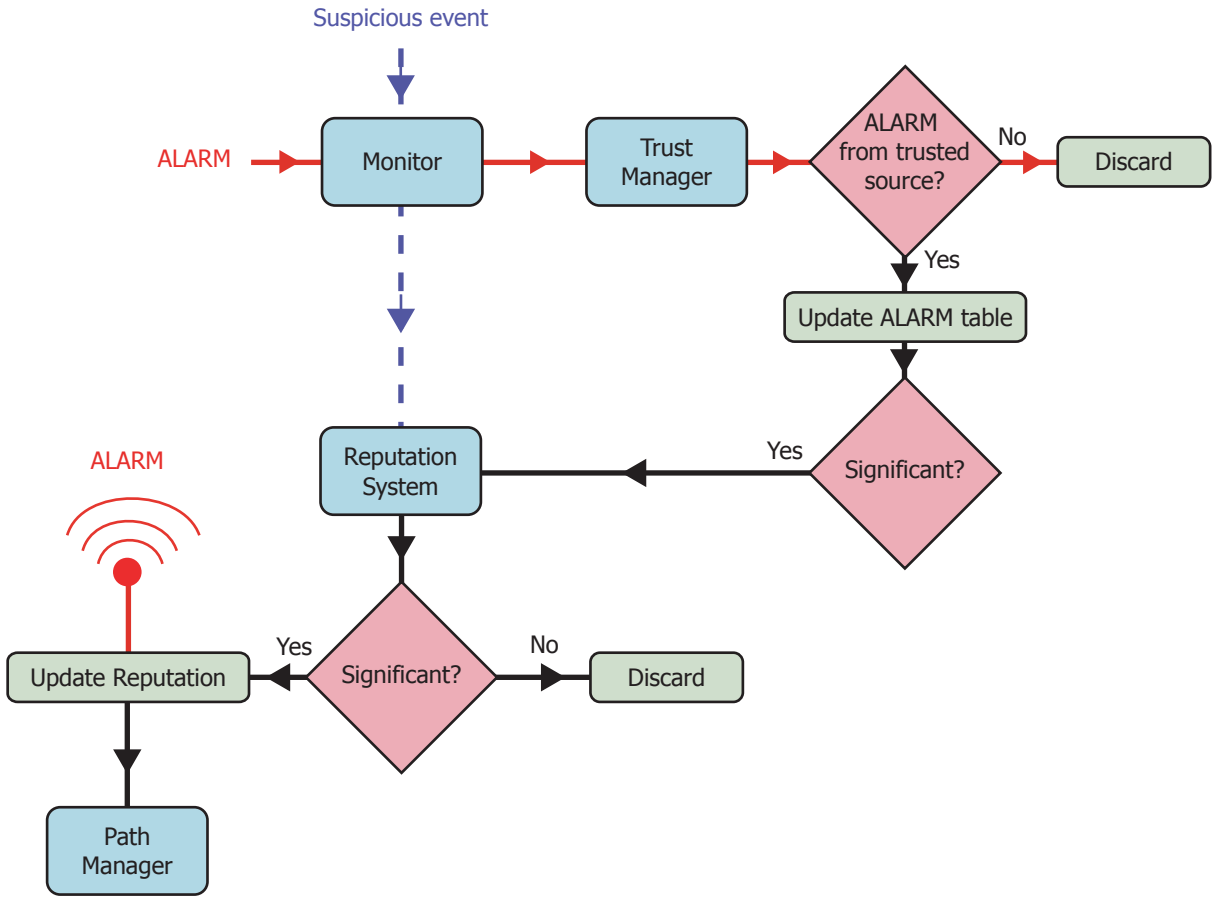

Figure 1: The interaction between the components of CONFIDANT.

another function. Thus, global reputation values do not reveal the importance placed on different services by different nodes.

CORE and CONFIDANT calculate reputation values in different ways. CORE reputation values range from positive (+1), through null (0), to negative (-1). CONFIDANT, however, only uses negative values. There are advantages in both approaches. Having a positive to negative range allows good behaviour to be rewarded and bad behaviour to be punished. By placing more weight on past behaviour, the CORE scheme is tolerant of sporadically bad behaviour. Thus, if a node with a good reputation is only temporarily unable to perform some function, e.g. because of environmental conditions, it will not be punished severely. However, such an approach is vulnerable to an attack where a node can build up a good reputation before behaving maliciously for a period. The better the reputation the malicious node can build up, the more time the node will have in which it can fail to meet its obligations while its reputation is still positive. There is also an issue with the way CORE places emphasis on past behaviour to eliminate sporadic activity, since inconsistent behaviour should probably be penalised, since it may be very damaging. Thus, while having a positive to negative range of reputations has advantages, placing weight on past experience may lead to undesirable effects.

This is an inherent fault with the reputation value used in all systems, i.e. an assumption that past behaviour is assumed to be indicative of future behaviour. Attacks involving 'building up credit' before behaving selfishly have less effect in CONFIDANT, as good behaviour is not rewarded, so all nodes are always under suspicion of bad behaviour. However, this makes CONFIDANT less tolerant of failed nodes, which may be exhibiting failed behaviour due, for example, to loss of power. In CONFIDANT, failed nodes recover through timeouts, when their entries in the blacklists expire and are deleted. This sudden leap back to a well behaved status will, however, enable a malicious node to repeat an attack.

Finally there is an issue regarding the possible negative effects of a reputation system on well-behaved nodes. A node which has built up a good reputation for a service may become a bottleneck in providing that service. This is an unwanted consequence of good reputation; the good node may even wish to decrease its reputation by behaving badly to prevent its resources being over-used. This would clearly be undesirable. This can be mitigated by placing more weight on unwanted behaviour than positive behaviour in the reputation calculation, so that unwanted behaviour will change the reputation value more than any desired behaviour.

4.1.2 Reputation and Identities

A reputation value should only be associated with one node. Thus node identification is a vitally important issue, and one which has not yet been adequately addressed in the literature. In order for reputations to be reliable, each node in the ad hoc network has to be bound to one identifier, so that a node cannot change identity in order to get a new reputation value.

The need for persistent identities will also depend on how the system treats new nodes. Each reputation system has a learning period, as the network will not know how a new node will behave. Friedman and Resnick [3] present work on discouraging new participants from malicious behaviour by assigning them the lowest possible reputation value. They argue that this promotes identity persistence in circumstances where there exists the ability to change identities easily. If a new node is allocated a low reputation rating initially, it has to perform positive work to gain a good reputation, in order that other nodes do not refuse service to the new node. Thus, a node which changes its identity will automatically lose any good reputation it has achieved.

Unfortunately, the use of this mechanism may not be feasible in an ad hoc network, where instantaneous connection is required. For example, in the accident data reporting example, the delays introduced by such a scheme would be unacceptable. Thus both CONFIDANT and CORE allocate neutral reputations to new nodes. Assigning a null value is a reasonable approach only if identity persistence can be encouraged by some other means.

4.2 Detecting Unwanted Behaviour

In order for reputation values to be trusted, nodes will need a reliable means of detecting good or bad behaviour. The process that is chosen will have to be dependable, where the objective is to eliminate subjectivity in the calculation of the reputation value. This is problematic, as bad behaviour by one node may be defined as good behaviour by another node. In the CORE scheme, for example, an access point may refuse to forward a packet for a node because it has not presented the correct authentication credentials. If the node has a good reputation in the view of another observing node, the observing node will define this action as selfish misbehaviour, and will decrease the reputation of the access point.

This example also highlights how detecting unwanted behaviour is the responsibility of both the nodes directly involved and the surrounding neighbour nodes. Both CONFIDANT and CORE rely on promiscuous observation for monitoring operations such as packet forwarding. However, passive observation presents several weaknesses when used within a wireless ad hoc environment. Some of these have been identified by Marti et al. [4].

Interoperability issues mean that in order to perform promiscuous observation the network hardware may have to be able to listen and construct packets over different radio transmission technologies. This may not be possible for resource constrained devices. Also, data packet collisions may occur when two nodes try to transmit at the same time. Thus, a node 'listening' for its packet to be forwarded may never 'hear' the packet, because a collision could have occurred while its neighbour node was forwarding the packet. The node receives no passive acknowledgement (i.e. does not see the packet being forwarded) and hence decreases the reputation of its neighbour node, even though the neighbour node successfully forwarded the packet. It may also be the case that the forwarding node moves out of range of the originator node to deliver the packet, in which case the 'listening' node will never receive the passive acknowledgement. An active acknowledgement mechanism would help to alleviate these problems but will add more overhead, and it will not solve the passive acknowledgement problems for neighbour nodes which are not directly involved in a communication between two of their other neighbours.

Reliance on promiscuous observation also limits a scheme's ability to function with other technologies such as directional antennae and wired networks, where promiscuous observa-

tion cannot occur since the transmission of the forwarded packet will never be received by a 'listening' node with just a directional antenna. This would mean that nodes would need an omni-directional antenna just for promiscuous observation, which would typically be uneconomic.

Detection can only occur if there is an event to detect. Thus the reputation value is only useful when nodes communicate more than once. The reputation value becomes more reliable as more communication takes place. CONFIDANT assumes that all nodes send the same amount of packets at a constant rate[4], i.e. each node performs the same number of services for each other node. Thus there is nothing to prevent a sleep deprivation torture attack where one node constantly sends packets to another node to forward. However, it is debatable as to whether or not this attack should lie within the scope of reputation systems.

Another issue is detecting 'selective misbehaviour', where a node deliberately seeks to minimise its use for routing packets, i.e. it tries to ensure that it is not included in any routes, whilst behaving correctly in other respects. By never dropping packets which it ought to forward, it should never receive a bad reputation (although it might also fail to receive a positive reputation). How this might be done will depend on the routing protocol, but it could for example be achieved by sending false or null responses to route request messages.

4.3 Analysis of a Distributed Reputation System

Both CORE and CONFIDANT distribute reputation values to create a network global view of a node's reputation. Any bad behaviour directly experienced by a node will be relayed to the whole network, so that bad behaviour is discouraged more than if reputation remained local knowledge. CORE requires nodes to send positive messages about cooperative nodes, and messages about badly rated nodes who are requesting services from good nodes. CONFIDANT sends ALARM messages, reporting any bad behaviour which is experienced or observed.

Many of the problems faced by distributed reputation schemes are the same as those faced by any other distributed scheme. For example, reputation messages could be modified or replayed. Moreover reputation messages may themselves be accidentally lost. As a result, there is a strong likelihood that serious inconsistencies will arise within a community as to the reputation values for nodes in the network. If these inconsistencies can be exploited by other nodes, then this is a serious vulnerability. Another important issue is the volume of additional messages which may be needed to support the distributed system. As bandwidth may be very limited, the priority may be on using the available bandwidth for emergency data rather than for reputation information. Thus, it seems that CORE and CONFIDENT will be better suited to ad hoc networks which use a proactive routing protocol, so that reputation information can be piggybacked on periodic route update information. If a global reputation system was to be used with an ad hoc network based on a reactive routing protocol, where information about unused routes in not stored, then serious inefficiencies will arise. The issue of storing information about every node in the network is also important. To use CORE and CONFIDANT, nodes will need a reasonable amount of storage space relative to the network size, as the systems require them to store information relating to all nodes.

Finally we mention three types of behaviour which can give rise to threats when reputation values are distributed throughout an ad hoc network:

- Advertising a false high rating about another node,
- Advertising a false low rating about another node, and
- Negative discrimination, where a node refuses services to only some nodes; this can be random or targeted at certain nodes.

4.3.1 Advertising false high reputation values

Advertising a false high rating is an attack which can be achieved in the CORE scheme, where malicious nodes can send positive rating messages to boost the reputation of nodes which have bad reputations. As noted in [5], the malicious node itself has not gained any direct advantage, but if it is in collusion with the badly rated nodes it has helped then this indirect attack is potentially a genuine threat. An example of a real world scenario is a user with a Personal Area Network (PAN), who does not want all of the devices in the PAN providing services in the ad hoc network. Thus the user involves just one device, which relays positive information about the other devices in the PAN, so that they all remain trusted and so that they can use the ad hoc network to request services. CORE does give some protection, in that subjective reputation has more weight than indirect reputation in the reputation calculations.

The authors of CORE deliberately designed their scheme to only allow positive reputation messages, to prevent malicious nodes advertising unfairly low reputation ratings. However, CORE still uses denial-of-service messages to report nodes with bad reputations trying to access services from good nodes. As mentioned above, neighbour nodes must check that the reputation values contained within any received denial-of-service messages correspond to the values stored in their own reputation tables. If the values do not match, the peer nodes decrease the reputation of the node which sent the denial of service message. However, the frequency of mismatches may be high, due to nodes moving in and out of range of one another, and the different assignment of importance to different functions, as outlined above.

4.3.2 Advertising false low reputation values

Transmitting false low reputation messages about a node is a denial of service attack, where decreasing a node's reputation will result in good nodes refusing it service. In both this and the previous attack, assurances will have to be given to a node of the reliability of the reputation value, so the node can decide the level of the risk before deciding whether it is willing to accept the reputation value. CONFIDANT tries to achieve this through a distributed trust mechanism, which is used to filter false ALARM messages from true ones. However, there are still some important issues.

CONFIDANT relies on trusting nodes to report bad behaviour. Two nodes with the same trust level could report ALARMS about each other. It is not clear what will happen in this event, but a reasonable conclusion will be that the first node who manages to distribute enough ALARM messages to exceed the reputation threshold will be trusted

[4]CORE does not explicitly assume this, but it still suffers from the same problems.

first. This attack would be easier for a node with a higher level of trust, which could reduce the reputation of nodes which have been assigned less trust. In view of this, the reputation system may motivate nodes to retaliate against ALARM messages, by targeting the node which originated the ALARM messages. Retaliation can also be achieved by conventional denial-of-service attacks such as sleep deprivation torture. It could be the case that it is more advantageous for a node to directly attack a node who is trying to send ALARM messages, than to let the ALARM message be processed so that its reputation is decreased.

Thus, the CONFIDANT scheme depends on how trust is placed and managed. Trust could be defined at a manufacturer level, or on an ownership level, for example. However, it may be infeasible to use trust assigned in this way to detect bad behaviour in an ad hoc network, which is vulnerable to both malicious and non-malicious failures.

4.3.3 Discrimination

Negative discrimination is a very important attack to consider. In a distributed system, a node could keep its overall network reputation high by cooperating with more than a certain of percentage of the nodes. This would suit attacks such as partial dropping of packets by selfish nodes.

There is also an issue with how the scheme does not discriminate between failed and selfish nodes, and between badly failed and malicious nodes. For example, failed nodes may fail because of factors outside of their control such as not having enough resource to perform a service. A selfish node will not perform the service, even though it has the resources to do so. CORE treats both types of nodes in the same way. The problem is that there is no way for the failed node to recover its reputation when it recovers from its own problems, as no nodes will request services from the failed node to enable it to perform well to increase its rating.

5. A SIMPLE REPUTATION MECHANISM

Dellarocus [2] and Friedman and Resnick [3] present a solution to some of the above problems in the context of a conventional reputation system; these solutions use 'controlled anonymity' to mitigate the attacks. A malicious node cannot send false reputation messages about another node if it does not know the node's true identity. While negative discrimination is still possible, it cannot be targeted at certain nodes. However, controlled anonymity requires some form of central control which is not feasible in vehicular ad hoc networks.

From the analysis in this paper, we can conclude that the reputation value should relate to exactly one function. The most reliable and quickest reputation values are those which are directly derived from personal experience. Both good and bad behaviour should influence the reputation value, where greater weight is placed on bad behaviour.

It makes sense to restrict reputation systems to just local calculations, due to the difficulties outlined above in synchronising reputation data. Also, reputation information is only going to be of use to the nodes in the area surrounding the badly behaved node.

Thus the following is a localised simple reputation mechanism is proposed. A node maintains a reputation table which consists of entries for every neighbour and their reputation for performing a certain function. The reputation value *rep*, is initially set to the variable *startrep*.

When a node requests a service from a neighbour, it gives the neighbour x opportunities to respond, where initially x is equal to *rep*. If the response is positive, x is increased by *changerep*. While x is positive, the value of x should be returned to the initial starting value after a *timeout* period, so that the reputation has to be earned again[5]. After a certain number of consecutive *timeout* periods where no negative behaviour has occurred, the *rep* value should be increased by *changerep*.

Where there is no response or the response is negative, x is decreased by $2\times$*changerep*. The node should keep trying until x reaches zero, when the corresponding *rep* value is decreased by $2\times$*changerep*. In this event, the node should look to request the service from a different node. If later on, the node wishes to try and request the service from the same neighbour again, it performs the same algorithm, where the *rep* value is less and thus the number x of opportunities is now less, i.e. the neighbour is given less chances. The node should perform exponential backoff to allow the neighbour to recover from any temporary problems.

Neighbour nodes should be given some chance of recovery. Thus, if a node has no other option but to try a selfish node, the node can just request the service with an initial x value of 1. This ,along with a decreasing *rep* value, results in less resources being wasted on a neighbour which is selfish or failed. Also, to discourage unwanted behaviour, service requests from nodes with reputation values below a threshold should be ignored.

The variables *startrep* and *changerep* used in this algorithm will have to be made subject to simulation tests to discover the optimum values. It is likely that they will depend on several factors including mobility and frequency of failure.

The reputation value can be used as an indication of the neighbour node's ability to perform a service. Thus when faced with several neighbours who offer the same service, a node can use the reputation value as a metric to influence its decision.

6. CONCLUSIONS

Mobile ad hoc networks have a number of significant security issues, especially those relating to inter-node routing. Various types of attack exist from nodes internal to the network. Reputation systems are used to establish trust and encourage trustworthy behaviour. CORE and CONFIDENT are two distributed reputation systems which have been proposed to mitigate the effects of internal threats in ad hoc networks. Unfortunately, reputation systems have inherent problems in the way the reputation value is defined and calculated, the detection of disreputable behaviour, and the coordinated distribution of reputation information.

Ongoing research involving detailed simulations of specific proposals for reputation schemes will help decide the way forward in this area. However, the advantages of reputation systems can still be enjoyed in a localised mechanism, where both positive and negative behaviour is included in the reputation calculation. Negative behaviour should be given greater weight than positive behaviour, so that when it occurs it severely affects the reputation value.

[5]This makes it more difficult for a malicious attacker to build up a good reputation to attack for a sustained period of time.

7. REFERENCES

[1] S. Buchegger and J.-Y. Le Boudec. Performance analysis of the CONFIDANT protocol (cooperation of nodes: Fairness in dynamic ad-hoc networks). In J. Hubaux, J. J. Garcia-Luna-Aceves, and D. Johnson, editors, *Proceedings of The Third ACM International Symposium on Mobile Ad Hoc Networking and Computing, 9-11 June, 2002, Lausanne, Switzerland*, pages 226–236. ACM Press, 2002.

[2] C. Dellarocas. Immunising online reputation reporting systems against unfair ratings and discriminatory behaviour. In A. Jhingran, J. Mason, and D. Tygar, editors, *Proceedings of the 2nd ACM conference on Electronic commerce, 17-20 October, 2000, Minneapolis, Minnesota, USA*, pages 150–157. ACM Press, 2000.

[3] E. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.

[4] S. Marti, T. J. Giuli, K. Lai, and M. Baker. Mitigating routing misbehavior in mobile ad hoc networks. In R. Pickholtz, S. Das, R. Caceres, and J. J. Garcia-Luna-Aceves, editors, *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking, August 6-11, 2000, Boston, Massachusetts, USA*, pages 255–265. ACM Press, 2000.

[5] P. Michiardi and R. Molva. CORE: A collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. In B. Jerman-Blazic and T. Klobucar, editors, *Communications and Multimedia Security, IFIP TC6/TC11 Sixth Joint Working Conference on Communications and Multimedia Security, September 26-27, 2002, Portoroz, Slovenia*, volume 228 of *IFIP Conference Proceedings*, pages 107–121. Kluwer Academic, 2002.

[6] L. Mui, M. Mohtashemi, and A. Halberstadt. Notions of reputation in multi-agents systems: a review. In M. Gini, T. Ishida, C. Castelfranchi, and W. Johnson, editors, *Proceedings of the first international joint conference on Autonomous agents and multiagent systems, July 15–19, 2002, Bologna, Italy*, pages 280–287. ACM Press, 2002.

[7] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. In M. Baye, editor, *Advances in Applied Microeconomics: The Economics of the Internet and E-Commerce*, volume 11, pages 127–157. Elsevier Science Ltd., November 2002.

[8] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.

[9] P. Yau and C. J. Mitchell. Security vulnerabilities in ad hoc networks. In *The Seventh International Symposium on Communication Theory and Applications, July 13–18, 2003, Ambleside, Lake District, UK*, pages 99–104. HW Communications Ltd, July 2003.