Limits on Revocation in VANETs

Bisheng Liu[1], Jerry T. Chiang[2], and Yih-Chun Hu[2]

[1] School of Computer Science
Fudan University, Shanghai, China
bsliu@fudan.edu.cn
[2] Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, Illinois, U.S.A.
{chiang2, yihchun}@illinois.edu

Abstract. We examine the limitations on revocation approaches in VANETs, including local revocation and global revocation. Local revocation schemes often use a local vote to identify and revoke an attacker. However, such votes often require that not only a majority of local nodes are honest, but that they also are able to detect the attack. We argue that these requirements may not be practical, particularly in the early stage of VANET deployment. Another local revocation approach, RevoGame [1], uses game theory to mitigate misbehavior in the VANET; however, we argue RevoGame does not correctly identify the players in the revocation game. We also analyze the limits of global revocation based on the misbehvaior accusations made and evidence gathered by each vehicle. Our analysis shows that no algorithm that uses only the accusation graph can identify attackers without false positives and false negatives.

Keywords: vehicular ad hoc networks, certificate revocation, graph theory

1 Introduction

According to the Center for Disease Control and Prevention, motor vehicle accidents are the leading cause of death among Americans between 1 and 34; the National Highway Traffic Safety Administration (NHTSA) reports that 37,261 Americans died in traffic accidents in 2008 [2]. Improving traffic safety is thus of the utmost importance. Moreover, the importance of improving traffic safety is not particular to the United States, but global. In China, for example, 73,484 died and 304,919 were injured in traffic accidents in 2008 [3]. Many of these accidents could have been avoided or made less severe if the driver were given advance warnings. Researchers have proposed Vehicular Ad Hoc Networks (VANETs), in which vehicles could form a multi-hop network and disseminate safety messages. In order to provide a reliable safety-of-life service, VANETs must be resilient to misbehaviors as well as universal service.

In order to preclude any attacks on the network from a malicious party, such as injecting misleading and false messages, a vehicle should carry with it a credential in order to join a VANET. That is, an entity must be granted explicit access to a VANET before it can communicate to other members of such network. To prevent any leak of privacy, a trusted third party known as the Certificate Authority (CA) is generally assumed to manage the identities, credentials, and cryptographic keys of all nodes in a VANET; this structure can mirror that used for issuing vehicle identification numbers (VINs).

If a node in the network can be shown to be misbehaving, be it by a set of other nodes or by the CA, its credential can be revoked so that a legitimate node will reject future communications from the misbehaving node. In *local revocation*, when a node is found to be misbehaving by a set of other nodes, the detecting set may cooperatively revoke the credential of the misbehaving node from their neighborhood. In *global revocation*, a CA finds a node to be misbehaving, possibly using input from other nodes, and revokes the misbehaving node's credential from the entire network.

In local revocation, a set of nodes can revoke an accused node by *voting* (e.g., [6-11]). Once a node is found to be malicious by a vote, participating nodes can reject future messages from the revoked node; this rejection is typically limited to the revoked key; any permanent revocation is done in conjunction with a global revocation scheme. This is because of the privacy concerns inherent in allowing a user to correlate two different keys of the same revoked node.

If a CA determines that a node is misbehaving, the CA can simply broadcast a revocation list containing the credential of that node to all nodes in the VANET [4, 5]. Suppose the CA can not observe all of the network all the time, it must then rely on accusations made by individual VANET nodes in order to decide whether or not a node should be revoked.

One of the most difficult challenges of revocation in VANETs is how to balance security and privacy. On one hand, most proposed revocation mechanisms try to exclude malicious attackers by exposing their identities and revoking their credentials; on the other hand, some drivers may value their privacy above the added safety of VANETs and will never willingly adopt a system that relinquishes their anonymity. If a vehicle is assigned only one certificate, then the identity of that vehicle can be inferred from its certificate; researchers have thus proposed protocols wherein each vehicle uses randomly changing certificates, known as pseudonyms [12], to sign messages. If the VANET uses a pseudonym scheme, revoking a node is much more difficult, since the revoked node can easily switch to another pseudonym when one of his pseudonym is revoked.

Existing local revocation schemes often assume an honest local plurality; that is, they assume that more nodes will vote against an attacker than vote in favor of the attacker. This represents a plurality because some nodes may have insufficient information to decide. The assumption of an honest local plurality faces four potential problems: the potentially limited set of nodes that can detect an attacker, the roadway may be so sparse that a few attackers can comprise a majority, the possibility of an attacker using several pseudonyms from the same vehicle, and the possibility that an attacker compromises keys from several vehicles. Certain researchers have thus

suggested that vehicles should employ secured hardware, such as a tamper-proof device to store security credentials [18, 32]. However, tamper-proof devices can significantly increase the cost of deployment of VANETs.

Our main contributions in this paper are:

- Our analysis challenges the current work in local revocation, because of the difficulties of achieving an honest plurality, especially when attacks are hard to detect. Furthermore, we show that sparse environments, such as those present when VANETs are first deployed, pose substantial challenges to local revocation.
- We show that a proposed game-theoretic approach [1] misidentifies the players in the revocation game and argue that the actual players are the manufacturers of vehicles, rather than individual drivers.
- We prove that any global revocation scheme that uses only accusations between nodes must either have false positives or false negatives.

The remainder of this paper is organized as follows. In Section 2, we begin with related work, present our assumptions and attacker models, and give the problem statement. In Section 3, we analyze local revocation, including limits of local revocation and the game theoretic local revocation problem. In Section 4, we use a graph theoretic model to show fundamental limits of global revocation. Finally we conclude our work in Section 5.

2 Background

2.1 Related Work

Previous work considers security and privacy in VANETs. Parno and Perrig [13] provide an overview of the challenges. Other work considers general security requirements and architectures for VANETs [14-16].

Prior research also studies the problem of authentication in VANETs [17-18]. Most such work adopts a *public key infrastructure* (PKI), in which a *certificate authority* (CA) issues one or more *certificates* to each vehicle, attesting to that vehicle's legitimate participation in the VANET [19]. Because the CA is able to introduce arbitrary nodes into the network, the CA must be trusted; for example, the CA can be a government authority. Most work assumes that each vehicle is equipped with many certificates to reduce privacy loss [12].

Other work focuses on the distribution of revocation information in VANETs. Raya et al. [11] present three mechanisms to distribute revocation information to all vehicles. First, they propose using Bloom filters to compress CRLs. Second, to revoke a vehicle entirely, the CA generates a revocation message that tells the tamper-proof device of the vehicle to stop all security functions. Third, if the CA is unavailable, the neighbors of the attacker vehicle could temporarily revoke the attacker as long as the number of accusing users exceeds a threshold. Papadimitratos et al. [21] propose breaking the *Certificate Revocation List* (CRL) into different pieces, then transmitting

these pieces using Fountain or Erasure codes, so that a vehicle can reconstruct the CRL after receiving a certain number of pieces. Laberteaux et al. [22] propose a mechanism for quickly distributing CRL through car-to-car (C2C) communication. Their analysis shows that very few Road-Side Units (RSUs) are required to distribute CRLs. Haas et al. [30] extend this mechanism by applying several optimizations. They suggest tying together all of a node's certificate identifiers through the use of a single key, so that a node's many certificates can be represented by a single value on the CRL. In addition, they exchange CRL updates, rather than the whole CRL, to limit network overhead. Our research differs from these approaches because we discuss the *decision making process of revocation*, rather than *how to distribute the revocation information*.

The work most similar to ours is in the literature of general mobile ad hoc networks, which often uses threshold cryptography [23]. Arboit et al. [7] propose that a node should be revoked when the sum of the weighted accusations against him is equal to or greater than a configurable threshold. The weight of each accusation depends on the trustworthiness of the accusing node. Similar ideas have been extensively explored [6, 8, 10]. Raya et al. extend this idea in vehicular ad hoc networks in a voting scheme, LEAVE [11]. In LEAVE, each vehicle that detects an attacker broadcasts warning messages. Once a vehicle receives enough warning messages about a suspicious node, the vehicle adds the suspect‘s identifier to a local blacklist and propagates “disregard” messages to its neighborhood together with the associated supporting signatures contained in the received warning messages. A vehicle receiving the disregard message can choose to ignore all future messages from the accused attacker. Unlike other accusation approaches, LEAVE temporarily revokes a user but leaves long-term revocation decisions to the CA, which may use evidence LEAVE provides. To be secure, most of these voting based approaches require an honest majority. However, our research questions the general applicability of that assumption in vehicular networks.

Suicide [9] is another proposed approach aimed at preventing an attacker from falsely voting against legitimate nodes. In suicide, a node can revoke another node by invalidating the credentials of both the accuser and the accused. Moore et al. propose a modified extension in VANETs[25]. Suicide based approaches only work well when certificates are a scarce resource. However, when each node is loaded with many pseudonyms, neither the accuser nor the accused are substantially affected, since both of them can easily switch to another certificate.

Raya et al. use game theory to analyze vehicle actions [1]. A vehicle that detects an attacker has three possible strategies: vote against the attacker, abstain from voting, and commit suicide to revoke the attacker. The authors propose a game-theoretic revocation approach, RevoGame, which defines the best strategy for each individual vehicle. However, our analysis suggests that since manufacturers, rather than drivers, program these revocation protocols, manufacturers might program these protocols to minimize their own cost.

Other related research includes reputation systems, which can be divided into two categories: distributed and centralized reputation systems. Distributed reputation systems require each node to continuously monitor their neighbors [26]. However,

due to high-speed mobility in VANETs, the interaction time among vehicles can be very short and repeated interactions may be rare, making distributed reputation systems difficult to apply in VANETs. In a centralized reputation system, a central authority collects node observations and derives a reputation score for each node [27]. Centralized reputation systems are similar to global revocation approaches except that global revocation approaches make binary decisions about trustworthiness, whereas centralized reputation systems can represent a fine-grained scale of trustworthiness.

2.2 Assumptions

In this section, we present a few assumptions to help shape the problem statement.

VANETs will likely use a hierarchy of CAs, organized by geographic region and possibly by manufacturer. For clarity, our discussion assumes that there is only one CA, though our analysis can be generalized. The CA manages the identities, cryptographic keys, and certificates of all nodes in its region. We assume that every legitimate vehicle has a unique identity V, several pairs of private and public cryptographic keys, $(PrKV)_i$ and $(PuKV)_i$, and corresponding certificates, $Cert(PuKV)_i$. Only the CA can correlate the identity of a vehicle to any of its certificates, and only the CA can determine whether or not two certificates come from the same vehicle.

We assume that a vehicle signs each message the vehicle sends and verifies the signature contained in every received message, rejecting messages without a valid signature. As mentioned in Section 1, for privacy reasons [12], a vehicle uses each public key only once, and for a short duration. The duration is out of scope of this paper. We assume that an attacker with physical access to a vehicle has access to all of that vehicle's cryptographic keys. Though tamper-resistant hardware makes key extraction more difficult, highly-effective tamper-resistant hardware is very expensive, and a determined attacker can compromise less effective hardware.

We believe that detection is orthogonal to revocation. We assume that some vehicles have detection mechanisms. A vehicle with a detection mechanism will detect some fraction of attacks that take place within the detection range of that vehicle. Relatively few of these vehicles may be present, especially when VANETs are first deployed, considering the deployment of VANETs could be a long and slow process. We do not assume an honest majority, for reasons discussed in Section 3.

2.3 Adversary Model

In our model, the attacker is an *inside attacker* [18]; that is, it has access to legitimate credentials. Because we do not consider the problem of detection, the definition of misbehavior is beyond the scope of this paper, but we anticipate that it could be either intentional or unintentional.

We do consider specific attacks on the revocation protocol. For example, a *Sybil attacker* can simultaneously use multiple keys from the same vehicle, or use keys from multiple different vehicles. One or more attackers may deliberately accuse an honest node.

2.4 Problem Statement

This paper studies the inherent limitations of existing revocation approaches in Vehicular Ad Hoc Networks. Our analysis examines two properties: *who* is the decision maker of the revocation and *how* the revocation decision is made. The decision maker can be the CA or a group of nodes around the attacker. If the CA makes the revocation decision, is it possible to revoke a node *solely based on accusations from other nodes*? If the revocation decision is made among a group of nodes around the attacker, is it possible that *a good node gets revoked due to false accusations from attackers*?

In this paper, we discuss only the decision making process. We do not discuss the evidentiary basis for such accusations, nor do we discuss the means for distributing information about a revocation decision.

3 Local Revocation

In this section we discuss the limitations of local revocation approaches. In addition, we analyze the effectiveness of a previously proposed game-theoretic local revocation approach [1].

3.1 General Local Revocation

Some researchers consider local revocation to be the preferred method for revocation in vehicular networks [11, 25], because the CA may not always be available and the latency of global revocation approaches may be unacceptable. Local revocation usually works in a distributed manner; when a group of vehicles has deployed detection mechanisms or devices, each vehicle is able to detect a misbehaving node based on local detection in conjunction with observations made by other group members. These approaches typically involve a local voting process. Several local revocation approaches have been proposed. In most such protocols, once a fixed number or a fixed fraction of votes have been cast against a node, that node is revoked [6-11]. To increase the cost of false accusations, variations such as *suicide* [9, 25] have also been proposed, wherein a node can revoke an attacker by invalidating both its current certificate and the certificate of the attacker.

In protocols where a node is revoked once the number of accusations against that node reaches the threshold, an attacker that comprises a local majority is either able to avoid revocation or can arbitrarily revoke legitimate nodes. We show below a number of scenarios in which the attacker may comprise a local majority.

Many privacy-preserving protocols propose equipping each vehicle with many certificates. If an average driver uses his car 2 hours per day, and the protocol requires changing certificate every minute, each vehicle needs 43800 certificates every year [18]. Certificates can be either periodically loaded by a third party authority or derived from a master key shared between a CA and the node. For privacy reasons, only the CA can know the correlation between the certificates and the real identity of

the node, so a legitimate node cannot detect when an attacker uses two of its certificates simultaneously. Furthermore, if an attacker can simultaneously use several of its certificates, it will almost always have enough certificates to comprise a local majority.

None of the proposed privacy-preserving protocols yet prevent an attacker from simultaneously using two or more of its certificates, though such a protocol is conceivable. Even in this case, it is still not difficult for some attackers to legitimately get access to several cars to abuse the revocation system, such as a valet or an employee of a car rental company. Furthermore, an attacker that can remove certificates from used cars would almost always comprise a local majority.

Another limitation of local revocation is that some honest nodes may not be able to vote. As mentioned in Section 2.2, to detect an attack, a node must be equipped with appropriate detection devices or mechanisms, and such devices or mechanisms may have limited range. Under these circumstances, and considering the mobility inherent in vehicular networks, it is impractical to assume that every vehicle around the attacker has the ability to detect an attack. As a result, most honest nodes may be ineligible to vote due to lack of proper detection mechanisms, out of range, or not having yet deployed VANETs. On the other hand, attackers can accuse falsely at will without much cost. An attacker may be able to reach the threshold for evicting a good node, even at the cost of several certificates if *suicide* approaches are deployed.

3.2 Game-theoretic Revocation

Raya et al. [1] propose a game theoretic protocol of local revocation and suicide. Their analysis considers each vehicle to behave as a rational being, choosing to vote against an attacker or use suicide only when that behavior is in the vehicle's best interests. In particular, if an attack is mild, or if the vehicle thinks other nodes will vote against the attacker, the node might not vote against the attacker. Under these assumptions, the authors present a local revocation protocol.

Before we analyze their approach, we first briefly describe their model. They model revocation as a *finite dynamic (sequential) game*, where vehicles in a local area are the players. A vehicle detecting misbehavior has three options: *abstain* from the revocation procedure, which has no cost, to *vote* against the attacker, which has cost v, and to *sacrifice*, invalidating both the attacker's certificate as well as the vehicles own, which has cost 1. In their model, each vehicle's action depends on the actions of previous players and the anticipated actions of future players. They also consider the potentially increasing cost in each round if the attacker were not revoked in the previous round.

The authors define the set of players to be those that are able to detect the attacker, and those players act to minimize their individual costs. Thus, the *social cost* of this approach (that is, the cost incurred by all nodes, whether or not they are able to detect the attacker) is minimized only when it also simultaneously minimizes the personal costs incurred by each detector. The work relies on the assumption that each node has perfect knowledge of the number of remaining eligible voters in the neighborhood. However, this assumption may not be realistic, because a voter might not know the

detection capabilities of the neighboring nodes of an attacker. In addition, like previous local revocation protocols, the author's approach is also vulnerable to the *Sybil attack*.

One limitation of this approach is that many of the nodes around the attacker may be unable to detect the attacker, which could lead to substantial social costs that could have been mitigated if users were more altruistic. This limitation is an inherent result of the author's assumption that each vehicle makes a decision in its own best interest. However, we note that in real systems, most users do not program their own devices; rather, they leave the manufacturer-supplied software in place. This principle is well illustrated by locked cell phones, because the interests of manufacturers and users are not perfectly aligned in the mobile carrier industry. The optimal strategy for a cell phone user is to have an unlocked phone, because option has value. A manufacturer, on the other hand, locks its handsets so that the carriers are willing to subsidize the phone, resulting in higher sales. The iPhone platform is one for which we have accurate statistics regarding manufacturer restrictions and sales. In October 2009, about 15% of iPhones were jailbroken [20, 33], showing that the manufacturer's interest was much more heavily represented than the interest of individual users. Similarly, in most cases, a manufacturer, rather than an individual driver, will choose the vehicle's strategy.

If the players of the revocation game are manufacturers, then all manufacturers can cooperate to play as a single player if it is in their best interest. This approach will result in an optimal social good. Alternatively, each manufacturer can act as a single individual player in the revocation game. Vehicles from a same manufacturer might only share information among themselves to minimize the cost incurred by vehicles made by that manufacturer. We leave more sophisticated game models for future work.

4 Global Revocation

When a node is globally revoked, most nodes in the VANET will not have ever had contact with the revoked node, nor will most nodes have had first-hand experience with the misbehavior of the globally revoked node. Thus, nodes must establish trust about a set of remote accusations against a particular node in order to revoke it. For privacy and performance reasons, we cannot entrust vehicles with a list of these accusations. As a result, global revocation schemes typically use the trust between individual vehicles and the CA to bootstrap the trust needed for revocation. That is, we let a node upload its accusations against other nodes. The CA then identifies a set of misbehaving nodes and whether they should be revoked. Finally, the CA can broadcast a revocation list enumerating the credentials of the misbehaving nodes.

We *allow* (but do not require) the revocation mechanism to assume that a benign node would not accuse another benign node. This assumption requires a detection scheme with no false positives; that is, a benign node would never detect a benign neighboring node as a misbehaving node. The existence of such a strong local misbehavior detection scheme is not guaranteed; however, we will show that even if

such a local misbehavior detection scheme exists, a remote misbehavior detection algorithm based solely on accusations must still suffer from false positives (wrongfully revoking a benign node) or false negatives (wrongfully keeping an attacker).

If a benign node never accuses another benign node, each accusation the CA gets must be from one of the following scenarios:

- An attacker accuses a benign node
- An attacker accuses another attacker
- A benign node accuses an attacker

In other words, even though a benign node would only accuse an attacker, an attacker can accuse any node, attacker or benign.

4.1 Definitions

In this section we formally define several terms that will aid our analysis of global revocation.

Definition 1 (accusation graph): An *accusation graph* is an undirected graph $G = (V, E)$, where each vertex represents a VANET node and each edge in E denotes that one endpoint has accused the other endpoint of the link. We define the *distance* between two vertices in G as the length of the shortest path between them that is comprised entirely of edges in E. Disconnected nodes have an infinite distance by definition.

The CA constructs an accusation graph by drawing an edge between a pair of nodes when one of the two nodes accuses the other of misbehaving. Since an attacker can freely accuse other nodes, a node making accusation is as suspicious as the node being accused by it. Our model therefore uses an undirected graph since the two endpoints are both suspicious. If the CA maintains reputation information, that information can be used to assign a *weight* to each edge. To simplify our analysis, we choose to ignore edge weights; however, our conclusion extends readily to algorithms that use edge weights.

Definition 2 (radius of an accusation graph *G*): Let the *eccentricity* of a vertex v in G be the greatest finite distance between v and any other vertex. The *radius* of an accusation graph is then defined to be the minimum eccentricity of any vertex in G [31].

For example, as depicted in Fig. 1, the radius of the accusation graph is 2 and the eccentricity of vertex v is 2.

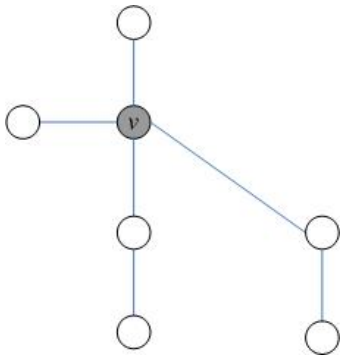

Fig. 1. An accusation graph G where the radius of G is 2

Definition 3 (trivial graph-based detection algorithm): A *graph-based detection algorithm* is a detection algorithm that returns a set of attackers based on the accusation graph and the accusation rules. The input of a graph-based detection algorithm is *(G, v)* where G is an accusation graph and *v is one of the nodes making accusations or being accused*. Assuming node v is a benign node, the output of a graph-based detection algorithm is a set of suspected attackers B_v which are *relative*[1] to node v. We then define a graph-based detection algorithm as trivial when the following holds: for any node $v \in V$, node $x \in B_v$ *iff* the edge $(x, v) \in E$. Here B_v denotes the set of attackers marked by the graph-based detection algorithm.

In other words, if we assume node v is a benign node, a trivial graph-based detection algorithm only marks the nodes that directly accuse node v or those that are directly accused by node v, that is, the direct neighbors of node v in the graph G.

Definition 4 (false positives and false negatives): A graph-based detection algorithm has *no false positives* when the following holds: for any accusation graph $G = (V, E)$, for any node $v \in V$, assuming node v is benign, then there are no benign nodes in B_v. A graph-based detection algorithm has *no false negatives* when the following holds: for any accusation graph $G = (V, E)$, for any node $v \in V$, assuming node v is benign, every attacker is in B_v.

4.2 Analysis

We now show the limits on global revocation based on accusation graphs.

Lemma 1: *Given input (G, v), if the radius of G is greater than 1, a graph-based detection algorithm with no false negatives would mark all the nodes other than v as attackers.*

Proof. It suffices to show that any algorithm that labels any node benign other than the input node must not be free of false positives.

[1] By *relative*, we mean that if a node n is in B_v, it means that when v is good, node n is bad.

For any accusation graph *G* with radius larger than 1, consider a graph-based detection algorithm that marks a node other than node *v* as benign. We present another valid outcome of the algorithm and show that the algorithm is not free of false negatives.

Since we do not allow any benign-benign accusation, if we assume *v* to be a benign node, our other possible benign nodes must be a distance of at least 2 away from *v*. The graph *G* has radius larger than 1, thus the set of *possible* benign nodes is not an empty set. An example accusation graph is shown in Fig. 2(a) where nodes *v* and *p* are labeled as benign.

However, since an attacker can accuse another attacker, node p can just as well be labeled as an attacker, as shown in Fig. 2(b), without violating our accusation assumptions. Since the CA, or any entity running a remote detection algorithm, has no way of verifying p to be benign with certainty, there is a non-zero probability that p is an attacker. Thus there are graphs for which the algorithm gives false negatives.

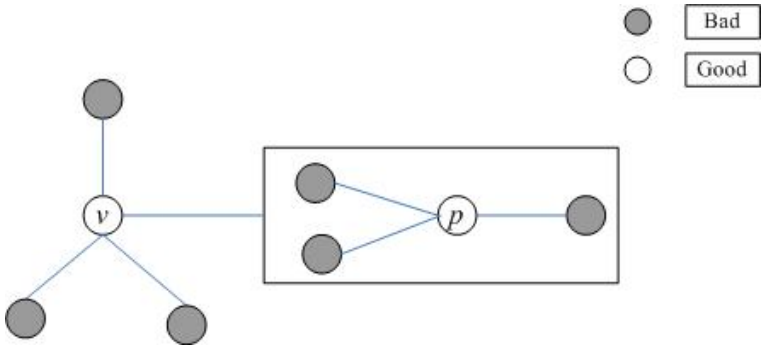

Fig. 2(a). An accusation graph where node *p* is labeled as good

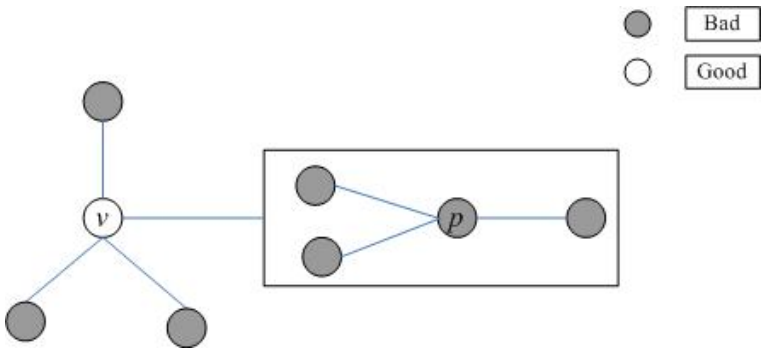

Fig. 2(b). Change the label of node *p* to bad

Lemma 2: *Given input (G, v), if the radius of G is greater than 1, a graph based algorithm with no false positives is a trivial graph based detection algorithm.*

Proof. It suffices to show that any algorithm that labels any node a distance of at least 2 away from v as an attacker must not be free of false positives.

For any accusation graph G with radius greater than 1, let there be a graph-based detection algorithm that labels a node that is at least a distance of 2 away from v as an attacker. We refer to this distant attacker as node q, as shown in Fig. 3(a).

We now change the label of q to a benign node, and change the label of all neighbors of q to attackers, as shown in Fig. 3(b). Since q is not a neighbor of v, our label changing scheme does not introduce any illegal benign-benign accusations. Again, since the CA cannot detect with absolute certainty which node is benign or misbehaving, there is a non-zero probability that a benign node is labeled as an attacker by the graph based detection algorithm.

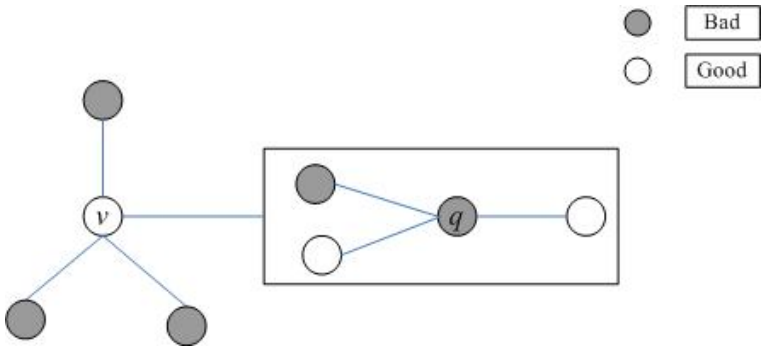

Fig. 3(a). An accusation graph where node q is labeled as bad

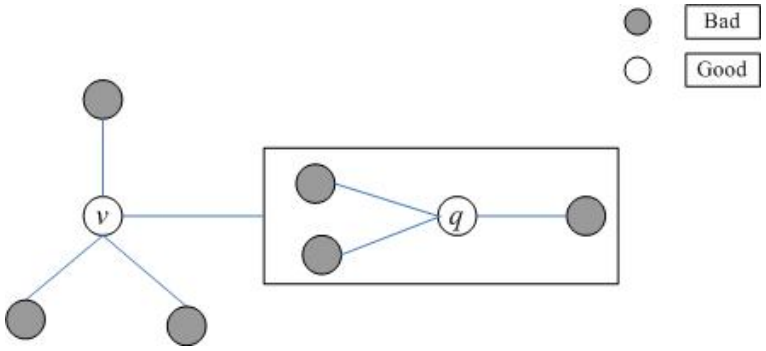

Fig. 3(b). Change the label of node q to good and the label of direct neighbors of node q to bad

Corollary 3: *Given input (G, v), if the radius of G is greater than 1, there is no graph-based detection algorithm that has no false positives and no false negatives.*

Proof. Suppose there is such an algorithm that has no false positives and no false negatives. According to Lemma 1, the algorithm should mark all the nodes other than node v as misbehaving. In addition, based on Lemma 2, the algorithm should only mark the direct neighbors of node v as misbehaving to guarantee no false negatives.

To meet both requirements, there should be no nodes other than node *v* and its direct neighbors in the accusation graph *G*, which contradicts the assumption that the radius of the input accusation graph is greater than 1.

5 Conclusion

In this paper, we analyze the limits on revocation in Vehicular Ad Hoc Networks (VANETs). Depending on the identity of the decision maker of the revocation process, revocation approaches can be generally classified into two categories: local revocation and global revocation. We examine the limits of both categories.

Local revocation does not rely on the availability of CA and can therefore quickly respond to attacks. Local revocation usually involves a local vote to revoke an attacker, which requires that most vehicles in the area around the attacker are not only honest, but also are able to detect the attack. Our analysis raises reasonable situations in which this assumption may not hold. We also challenge the choices made by a previous game-theoretic approach [1] and argue that players in the revocation game are likely to be vehicle manufacturers rather than individual users.

We also analyze the limits on global revocation. Global revocation relies on a CA to use the graph of accusations, together with any evidence gathered from vehicles, to make a revocation decision. Although global revocation introduces response delays, it gathers information from the entire network, meaning that the majority of nodes are quite likely to be honest. Unfortunately, we show that, when using only the accusation graph, no algorithm can mark all of the attackers without false positives and false negatives when the graph has radius greater than 1. This result shows that global revocation based on accusation graphs must necessarily accept false positives or false negatives. In light of these new results, extensive further work is required to build a revocation solution upon the existing work.

References

1. Raya, M., Manshaei, M. and Félegyházi, M.: Revocation games in ephemeral networks. In: 15th ACM conference on Computer and communications security, pp. 199-210. ACM New York, NY, USA (2008)
2. Fatality analysis reporting system web based encyclopedia, http://www-fars.nhtsa.dot.gov/Main/index.aspx
3. National traffic accident annual report, http://www.mps.gov.cn/n16/n85753/n85870/ index.html
4. Kocher, P.: On certificate revocation and validation. In: International Conference on Financial Cryptography. Lecture notes in computer science, 1465:172-177, (1998)
5. Zheng, P.: Tradeoffs in certificate revocation schemes. ACM SIGCOMM Computer Communication Review, 33(2):112, (2003)

6. Yi, S. and Kravets, R.: MOCA: Mobile certificate authority for wireless ad hoc networks. In: 2nd Annual PKI Research Workshop, pp. 65, (2003)
7. Arboit, G., Crépeau, C., Davis, C. and Maheswaran, M.: A localized certificate revocation scheme for mobile ad hoc networks. Ad Hoc Networks, 6(1):17-31, (2008)
8. Luo, H., Kong, J., Zerfos, P., Lu, S. and Zhang, L.: URSA: ubiquitous and robust access control for mobile ad hoc networks. IEEE/ACM Transactions on Networking (ToN), 12(6):1049-1063, (2004)
9. Moore, T., Clulow, J., Nagaraja, S. and Anderson, R.: New strategies for revocation in ad-hoc networks. In: 4th European Workshop on Security and Privacy in Ad hoc and Sensor Networks (ESAS). Lecture notes in computer science, 4572:232, (2007).
10. Chan, H., Gligor, V., Perrig, A. and Muralidharan, G.: On the distribution and revocation of cryptographic keys in sensor networks. IEEE Transactions on Dependable and Secure Computing:233-247, (2005)
11. Raya, M., Papadimitratos, P., Aad, I., Jungels, D. and Hubaux, J.: Eviction of misbehaving and faulty nodes in vehicular networks. IEEE Journal on Selected Areas in Communications, 25(8):1557, (2007)
12. Dotzer, F.: Privacy issues in vehicular ad hoc networks. In: Workshop on Privacy Enhancing Technologies (PET). Lecture notes in computer science, 3856:197-209, (2006)
13. Parno, B. and Perrig, A.: Challenges in securing vehicular networks. In: Workshop on Hot Topics in Networks (HotNets-IV). (2005)
14. Papadimitratos, P., Buttyan, L., Holczer, T., Schoch, E., Freudiger, J., Raya, M., Ma, Z., Kargl, F., Kung, A. and Hubaux, J.: Secure vehicular communication systems: design and architecture. IEEE Communications Magazine, 46(11):100-109, (2008)
15. Raya, M., Papadimitratos, P. and Hubaux, J.: Securing vehicular communications. IEEE Wireless Communications, 13(5):8, (2006)
16. Hubaux, J., apkun, S. and Luo, J.: The security and privacy of smart vehicles. IEEE Security & Privacy:49-55, (2004)
17. Hu, Y. and Laberteaux, K.: Strong VANET security on a budget. In: 4th Annual Conference on Embedded Security in Cars (ESCAR). (2006)
18. Raya, M. and Hubaux, J.: The security of vehicular ad hoc networks. In: 3rd ACM Workshop on the Security of Ad Hoc and Sensor Networks(SASN). pp. 21. ACM (2005)
19. Schneier, B.: Applied cryptography: protocols, algorithms, and source code in C. Wiley-India (2007)
20. http://www.pinchmedia.com/blog/piracy-in-the-app-store-from-360idev/
21. Papadimitratos, P., Mezzour, G. and Hubaux, J.: Certificate revocation list distribution in vehicular communication systems. In: 5th ACM International Workshop on VehiculAr Inter-NETworking (VANET). pp. 86-87. ACM (2008)
22. Laberteaux, K., Haas, J. and Hu, Y.: Security certificate revocation list distribution for vanet. In: 5th ACM International Workshop on VehiculAr Inter-NETworking (VANET). pp. 88-89. ACM (2008)

23. Desmedt, Y.: Threshold cryptography. EUR TRANS TELECOMMUN RELAT TECHNOL, 5(4):449-457, (1994)
24. Zhou, L. and Haas, Z.: Securing ad hoc networks. IEEE network, 13(6):24-30, (1999)
25. Moore, T., Raya, M., Clulow, J., Papadimitratos, P., Anderson, R. and Hubaux, J.: Fast exclusion of errant devices from vehicular networks. In: 5th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON). (2008)
26. Buchegger, S. and Le Boudec, J.: A robust reputation system for mobile ad-hoc networks. In: 2nd Workshop on the Economics of Peer-to-Peer Systems (P2PECON). (2004)
27. Resnick, P., Kuwabara, K., Zeckhauser, R. and Friedman, E.: Reputation systems. (2000)
28. Schneider, J., Kortuem, G., Jager, J., Fickas, S. and Segall, Z.: Disseminating trust information in wearable communities. Personal and Ubiquitous Computing, 4(4):245-248, (2000)
29. Specks, W., Matheus, K., Morich, R., Paulus, I., Menig, C., Lübke, A., Rech, B. and Audi, V.: Car-to-Car Communication–Market Introduction and Success Factors. In: 5th European Congress and Exhibition on Intelligent Transport Systems and Services. (2005)
30. Haas, J., Hu, Y. and Laberteaux, K.: Design and analysis of a lightweight certificate revocation mechanism for VANET. In 6th ACM International Workshop on Vehicular Ad Hoc Networks (VANET). pp. 89-98. ACM (2009)
31. Harary, F.: Graph Theory. Addison-Wesley, Reading, MA (1969)
32. Picconi, F., Ravi, N., Gruteser, M. and Iftode, L.: Probabilistic validation of aggregated data in vehicular ad-hoc networks. In: 3rd ACM International Workshop on VehiculAr Inter-NETworking (VANET). pp. 85. ACM (2006)
33. http://en.wikipedia.org/wiki/File:IPhone_sales_per_quarter.svg