Secure Incentives for Commercial Ad Dissemination in Vehicular Networks*

Suk-Bok Lee*, Gabriel Pan*, Joon-Sang Park‡, Mario Gerla*, Songwu Lu*
*Computer Science Department
University of California
Los Angeles, CA 90095
{sblee,pansh,gerla,slu}@cs.ucla.edu

‡Computer Science Department
Hongik University
121-791 Seoul, Korea
jsp@hongik.ac.kr

ABSTRACT

Vehicular ad hoc networks (VANETs) are envisioned to provide us with numerous interesting services in the near future. One of the most promising applications is the dissemination of commercial advertisements via car-to-car communication. However, due to non-cooperative behavior of selfish nodes or even malicious ones in the real-world scenario, such vehicular advertisement system cannot be realized unless proper incentives and security mechanisms are taken into consideration. This paper presents Signature-Seeking Drive (SSD), a secure incentive framework for commercial ad dissemination in VANETs. Unlike currently proposed incentive systems, SSD does not rely on tamper-proof hardware or game theoretic approaches, but leverages a PKI (Public Key Infrastructure) to provide secure incentives for cooperative nodes. With a set of ad dissemination designs proposed, we demonstrate that our SSD is robust in both incentive and security perspectives.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General—*Security and protection*

General Terms

Design, Security

Keywords

Vehicular ad hoc networks, Incentives, Security, Cooperation

1. INTRODUCTION

Vehicular ad hoc networks (VANETs) consist of smart vehicles on the road, and most of newly-manufactured vehicles are becoming smarter in no distant future. These vehicles are equipped with sensing devices, computing resources, and short-range radios (DSRC) [11] for communication with other nearby vehicles or with roadside infrastructure. They are envisioned to play an important role as "networks on the road", and provide numerous interesting services in the near future [26, 41, 44]. Among them, one of the most promising applications is to disseminate commercial advertisements via car-to-car communication, and some researchers have already visualized similar scenarios (e.g. Electronic Coupon Systems [22], FleaNet [24], and Digital Billboards [30]).

In such systems, an ad provider sends out commercial ads via vehicle-roadside communication [18], and nearby receiving vehicles start to disseminate those ads by forwarding them to other vehicles while moving (as in Figure 1). This vehicular dissemination system is very effective for advertisement, since a vehicle itself moves from place to place so that it can forward the ads whenever new vehicles move into its vicinity. In the ideal environment, each vehicle naively forwards those ads repeatedly for a certain time period. In reality, such a cooperative scenario, however, is not practical [27] in that some selfish users may not want to forward such ads for nothing. Forthrightly, even for naive users, why should they help forward those commercial ads for the benefit of the business companies? Also, as for security concerns, the more serious case is that, without proper security countermeasures, malicious nodes[1] who intentionally try to attack the network, for instance, may launch DoS (Denial-of-Service) attacks by sending out dummy ads propagating throughout the network. Thus, to deploy such vehicular ad system in real-world scenarios, proper incentives and security mechanisms should be taken into consideration.

To stimulate cooperation among those selfish nodes in mobile ad-hoc networks, several incentive schemes have been proposed so far. One possible approach is that, with the assumption on some degree of tamper-proof hardware [4, 5, 40] on each node, each relaying node can earn some virtual credit, which is managed and protected by the tamper-proof hardware, in order to motivate each individual node to participate in the network. Other incentive approaches make use of reputation-based schemes [2, 3, 21, 25, 27], which monitor neighboring nodes' traffic and keep track of the reputation of each other so that uncooperative nodes are eventually detected and isolated from the network. Also, several researchers have begun to investigate such non-cooperative communication scenarios within a game theory framework [13, 22, 42, 48]. By manipulating the parameters (e.g. the amount of gain per forwarding, the designation of charging subject, etc), those schemes encourage cooperative behavior among selfish nodes. However, as pointed out in [17], if poorly implemented in practice, these incentive schemes themselves have potential to backfire by offering an incentive to

*This work was supported in part by the US Army under MURI award W911NF-05-1-0246, the NSF under Grant No. 0520332, and the US Army Research Laboratory and the UK Ministry of Defence under Agreement No. W911NF-06-3-0001.

MobiHoc'07, September 9–14, 2007, Montréal, Québec, Canada.

[1]We use the terms "node", "user", and "vehicle" interchangeably.

cheat the system in order to gain further benefits. Thus, to prevent such undesirable flaws when designing an incentive scheme, careful investigation into blind points of the scheme is needed.

Compared with other aspects, security in VANETs has received little attention so far. As for the security building blocks for VANET, a Public Key Infrastructure (PKI) turns out to be the most suitable way for satisfying security requirements in car-to-car communication [32, 34]. Computing resources equipped in vehicles are capable to process digital signatures [47]. Also, unlike symmetric authentication mechanisms, asymmetric ones do not require a preliminary handshake, which is not acceptable most cases in inter-vehicle communication. With a PKI, each vehicle can have an *electronic license plate* [18] that provides its certified identity via a wireless link. Also, secure location verification schemes [6, 7, 23] for mobile nodes have been proposed, so we can verify a vehicle's location information as well as its identity. This kind of strong authentication provides valuable auditability for the authorities concerned, but it can breach drivers' privacy [8, 18, 32, 34]. Thus, balancing security with privacy is also an important issue in VANETs.

With these security concerns and uncooperative nature among users in mind, to realize such vehicular ad system, this paper presents Signature-Seeking Drive (SSD), a secure incentive framework for commercial ad dissemination in VANETs. Unlike currently proposed incentive systems, SSD does not rely on tamper-proof hardware[2] or game theoretic approaches, but leverages a PKI to provide secure incentives for cooperative nodes. We propose a set of ad dissemination designs and investigate their robustness in both incentive and security perspectives.

Inspired by a micro-payment scheme [20] and a charging/rewarding scheme [1], SSD employs the notion of *virtual cash* to charge and reward the provision of advertising service, as an incentive for users in the network. At a high level, the description of SSD is as follows. An ad-forwarding vehicle tries to obtain *receipts* from its ad-receiving neighbors. While driving its way, the vehicle may collect as many receipts as it forwards the ad. At the *virtual cashier* (e.g. gas stations), the vehicle can exchange those collected receipts with *virtual cash*, and the predefined amount of the cash is also reserved for each receipt-providing node. Paying in part by the virtual cash, drivers can, for example, gas up their vehicles. Later, an ad-providing company defrays the cost for the virtual cash induced by the ad, which stimulates cooperation among users.

We have evaluated our SSD through analysis and simulation experiments, and demonstrated its robustness in both incentive and security perspectives against various types of attacks.

The contribution of this paper is two-fold. First, we propose a potential and promising application scenario with a set of ad dissemination models. Second, we present a secure framework for commercial ad dissemination in VANETs, with both selfish users (incentives) and malicious users (security) into account.

The rest of this paper is organized as follows. Our system model is given in Section 2. We present SSD, our secure incentive framework in Section 3. An evaluation of SSD is given and discussed in Section 4. Finally, we present conclusions and future work in Section 5.

2. SYSTEM MODEL

In this section, we describe our vehicular network model and advertisement models, followed by the notations used in this paper.

[2]Each vehicle is likely to have a tamper-proof device. However, it is for protecting the secret information from attackers, not for an incentive mechanism itself.

2.1 Vehicular Network Model

In VANETs, each vehicle can communicate with other nearby vehicles or with fixed roadside infrastructure, to perform some useful applications [26, 33, 34, 36, 46] such as safety-related warning functions, traffic management, infotainment, payment services, etc. Considering high mobility nature of vehicles as well as their ad-hoc communication characteristics, VANETs are envisioned to substantiate the most promising model of mobile ad hoc networks [34].

Our vehicular network model assumes that each registered vehicle keeps its own certificate (i.e. public/private key pair issued by a Certificate Authority (CA)). There are two options for the designation of CA so far: governmental authorities or vehicle manufacturers. However, as indicated in [32, 35], an idea of the Department of Motor Vehicles (DMV) as a CA has several flaws mainly caused by administrative problems. As in [34], a notion of vehicle manufacturers as a CA is also disadvantageous in that the manufacturers are not as trustworthy as governmental institutions. In the real world, the distribution of key certificates to vehicles is still a challenging problem in VANETs, until new authority specialized for such operations appears.

In addition, to balance between strong authentication and drivers' privacy, each vehicle can have its temporary IDs (i.e. anonymous public/private key pairs). A vehicle can obtain its new anonymous key pair at the *reanonymizers* [32] stationed at regular intervals on the roadside, or a set of anonymous key pairs can be preloaded [34] in the vehicle. The certificates for such key pairs have short lifetimes so that third parties fail to track the real identities of vehicles, yet in some cases authorities are able to retrieve the actual entities from anonymous public keys. We will discuss this property related to our framework in Section 4.5.

2.2 Goal Model

2.2.1 Incentives

With the help of inter-vehicle communications, advertisers' goal is to disseminate their ads over the network, (1) within a certain time period, or (2) within a certain area. Example scenarios are as follows: (1) Electronic company S has new products to offer and has decided to launch an intensive advertising campaign for the first three months nationwide. (2) Drive-in fast food restaurant I has newly opened at Westwood area near UCLA campus and they want to advertise it mostly to drivers in the area.

Advertisers in all those types of scenarios probably want to use vehicular ad system to disseminate their ads, targeting a large number of potential customers inside the cars. However, from a viewpoint of vehicle users, those commercial ads are only for the benefit of the business companies and they are exploiting vehicle users' resources for their own profit. Users probably want some type of incentive to stimulate cooperation. Thus, the graceful compromise between these two sides is that advertisers pay for the incentives for users, in the sense that they pay charges for network resources, or advertising charges; in reality, for the commercial purpose, nothing is free.

2.2.2 Vehicular Authority

Every commercial ad here needs to get permission from the proper authorities before publicity. Here we assume to have an authority, which is specialized for the vehicular ad system, called a Vehicular Authority (VA). VA authorizes every vehicular advertisement and maintains the records of all the vehicular ad payment transactions. Thus, each vehicle is preloaded with the VA's public key as well as the CA's public key.

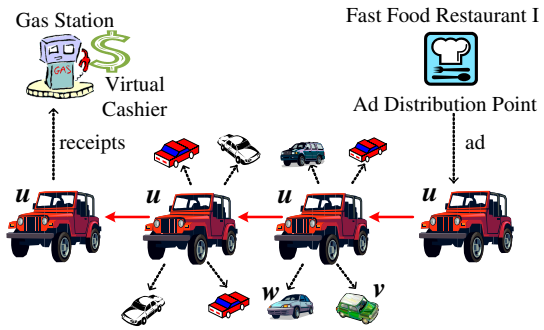

Figure 1: One-Level Ad Dissemination.

As in Figure 1, once vehicle u agrees to forward I's ad at *ad distribution point* (ADP), u continues to advertise it to any newly-encountered vehicle v for a certain time period. Ad-receiving vehicle v, in return, may provide a digitally-signed receipt for u. These receipts are exchangeable with virtual cash at the virtual cashier (e.g. gas station); the predefined amount of the cash is reserved for v, too. The virtual cashier sends all the transaction records to a VA. Then, a VA charges restaurant I such virtual cash induced by the ad, and I pays for what stimulates cooperation among users, as an incentive for both advertising node u and receipt-providing node v.

2.2.3 *Advertisement Level*

Our vehicular ad system has two options for disseminating ads, which are one-level and multi-level dissemination. In one-level dissemination, the only vehicles that receive an ad directly from ADP are allowed to continue forwarding the ad to others. In Figure 1, vehicle w can only receive the ad, and cannot reuse it unless w contacts with the ADP. This one-level dissemination is probably appropriate for local advertising such as scenario (2) above, because the advertiser's goal is not that every user in the area forwards the ad. Rather, he wants most users in the area to receive the ad, using a reasonable number of ad-forwarding vehicles sufficient to cover the area.

On the other hand, multi-level dissemination has no such restrictions. Figure 2 shows an example of multi-level dissemination. It allows vehicle v to reuse the ad from u, so v can make its virtual cash by forwarding this ad to other vehicle x. Then, x can also reuse and advertise it to others, which is similar to *pyramid selling*. This is more appropriate for the intensive advertising over a wide area such as scenario (1), which yet involves a large amount of advertising expenditure in that the advertising costs tend to increase with the number of ad-forwarding nodes. The detailed design of these two options will be described in Section 3.

2.3 Uncooperative Model

Like other mobile networks, VANETs have two types of uncooperative nodes [48]: selfish nodes and malicious nodes. They have specifically different goals in that selfish nodes seek to economically maximize their own profit, but malicious nodes try to attack the system with the intention of disrupting some part or even the whole network. A complete incentive system should deal with both selfish nodes and malicious ones.

We may encourage selfish nodes to participate in the network with an incentive model. However, malicious nodes still try to attack the weak point of the model. For example, malicious nodes may try to launch DoS attacks by disseminating dummy ads over the network, or colluding nodes may share their collected receipts and try to fabricate them to charge a large amount of virtual cash to advertising companies. Thus, to stimulate cooperation among users while preventing attackers from disrupting the system, we should pursue secure incentives.

We assume that any nodes cannot obtain a certificate for another entity from CA. In fact, preventing crafty attackers from cheating CAs is still a challenging problem in reality, and the management of CAs is always the problem of PKI.

Notations We use the following notations throughout the paper:

- u, v are principals, such as communicating nodes.
- C_u is u's certificate by CA.
- K_u^+ is u's public key.
- K_u^- is u's private key.
- $M1|M2$ is the concatenation of messages $M1$ and $M2$.
- $H(M)$ is hash of M (e.g. SHA-1 [12]).
- $\{M\}_{K_u^-}$ is u's digital signature[3] on M.
- AD_S is an advertisement by company S.

We also use $u \rightarrow v : M$ to denote that u sends message M to v, and $u \rightarrow * : M$ to denote that u broadcasts M.

3. SSD: SIGNATURE-SEEKING DRIVE

In this section, we present our secure incentive framework, Signature-Seeking Drive (SSD). We describe a set of advertising models in our framework.

3.1 One-level Advertisement

Approval for Advertisement As mentioned in Section 2.2, before publicity, every commercial vehicular ad has to be approved by a Vehicular Authority (VA) in charge of the vehicular ad system, and each vehicle is preloaded with the VA's public key as well as the CA's public key. Thus, restaurant I in scenario (2) gets its ad AD_I certified by a VA as follows:

$$I \longrightarrow VA: \quad C_I, AD_I, \{AD_I\}_{K_I^-}.$$
$$VA \longrightarrow I: \quad \{AD_I|I\}_{K_{VA}^-}.$$

where C_I is I's certificate. With I's digital signature on AD_I, VA authenticates I's identity. VA reviews it (e.g. the content of AD_I, I's previous records, etc) and then gives I an *ad-permit* $\{AD_I|I\}_{K_{VA}^-}$. The notion of an ad-permit prevents malicious users from disseminating unauthorized ads over the network.

Agreement with Ad Distribution Point When each vehicle u approaches I's Ad Distribution Point (ADP) (see Figure 1), the ADP contacts with u via vehicle-roadside communication as follows:

$$I \longrightarrow *: \quad C_I, AD_I, \{AD_I|I\}_{K_{VA}^-}.$$
$$u \longrightarrow I: \quad C_u, \{AD_I\}_{K_u^-}.$$
$$I \longrightarrow u: \quad \{AD_I|u\}_{K_I^-}.$$

where C_I and C_u are I's and u's certificates, respectively. Receiving I's ad AD_I, each vehicle u first checks whether AD_I has been approved by VA (i.e. ad-permit). If yes and AD_I is also for one-level advertisement[4], u may respond to this ad, with an incentive in mind, by sending back its certificate. Then, verifying u's identity, I provides u with a *voucher* $\{AD_I|u\}_{K_I^-}$ for u's exclusive use. Without this voucher, u cannot get its virtual cash at the

[3]Message M is hashed before being signed.

[4]Advertiser I indicates whether AD_I is for one-level advertisement, by specifying it in AD_I.

virtual cashier. Other vehicles fail to share this voucher, since it is tied with u's identity. Thus, the notion of a voucher limits the dissemination to one level so that the only vehicles contacting directly with I's ADP are allowed to make their virtual cash while forwarding the ad to others.

Advertisement Dissemination Hoping to make as much virtual cash as possible, u may advertise AD_I to any newly approaching vehicle v in behalf of I as follows:

$$u \longrightarrow *: \quad C_u, AD_I, I, \{\{AD_I|I\}_{K_{VA}^-}\}_{K_u^-}.$$
$$v \longrightarrow u: \quad C_v, \{AD_I|u\}_{K_v^-}.$$

where C_u and C_v are u's and v's certificates, respectively. Receiving AD_I from u, v first authenticates u' identity and checks whether this ad is certified by VA. If yes, v may respond to this ad, with an incentive in mind, by providing its certificate and a digitally-singed receipt $\{AD_I|u\}_{K_v^-}$ for u's exclusive use. Then, u stores them and may continue advertising in the hope of collecting as many receipts as possible. By simply reusing AD_I from u, vehicle v in this case fails to get virtual cash at the virtual cashier, since v does not have a legitimate voucher until it directly contacts with I's ADP.

Note that some selfish nodes even without vouchers still can reuse AD_I from u, advertise it, collect as many receipts as possible in advance, and then visit I's ADP later in order to obtain their proper vouchers. Our framework allows this kind of wily behavior in that, from a standpoint of advertiser I, the order does not actually matter (e.g. such nodes advertise AD_I first and get their vouchers later), as long as the one-level dissemination property is preserved at the virtual cashier. If those wily nodes, however, contact with I's ADP too late (e.g. after a certain time period, I may stop generating vouchers for AD_I), they fail to redeem their collected receipts without the proper vouchers at the virtual cashier.

Receipt Redemption Every certified ad has its own term of validity specified in it. Thus, vehicle u in Figure 1 has to redeem its collected receipts before the specified time of AD_I, otherwise the virtual casher refuses u the virtual cash due to the expired voucher. Before the expiration of AD_I, u may return its collected receipts to any nearby virtual cashier VC (e.g. gas station) as follows:

$$u \rightarrow VC: \quad C_u, AD_I, I, \{AD_I|I\}_{K_{VA}^-}, \{\{AD_I|u\}_{K_I^-}\}_{K_u^-}, R_u.$$
$$R_u = (C_v, \{AD_I|u\}_{K_v^-}), (C_w, \{AD_I|u\}_{K_w^-}), ...$$

where R_u is the receipts of AD_I collected by u. VC first checks 1) the due date of AD_I, 2) u's identity, and 3) the legitimacy of u's voucher. All those are verified with u's digital signature on its voucher $\{\{AD_I|u\}_{K_I^-}\}_{K_u^-}$. Then VC examines whether u has never redeemed u's voucher for AD_I at any other virtual cashier before, by inquiring of VA about u's previous record. Each VC is connected with a VA that maintains the records of all the vehicular ad payment transactions. Then VC verifies the legitimacy of each receipt from the "distinct" nodes in R_u (in this case, the receipts from v, w, etc). VC sends all of this data to a VA that keeps vehicular ad records.

Now u earns as much virtual cash as the number of the valid receipts in R_u.[5] The predefined amount of the cash is also given to each receipt-providing vehicle (e.g. v, w, etc.) in R_u. VC sends such virtual cash to each user's account. We envision each VC is connected with a virtual bank that maintains accounts for each vehicle. Then later a VA charges restaurant I such virtual cash induced by AD_I.

[5]Some advertiser S may set the maximum number of receipts that one vehicle can collect, by specifying it in AD_S.

Since a VA keeps all the records of vehicular ad payment transactions, u can redeem its voucher for AD_I only once. This prevents a selfish/malicious node from obtaining its virtual cash at different virtual cashiers by reusing the same receipts over and over. Thus, to maximize the profit, each vehicle u may submit its collected receipts to the virtual cashier as near the due date as possible.

3.2 Multi-level Advertisement

Some companies may want to advertise their products intensively over a wide area, as in scenario (1) in Section 2.2, without imposing restriction on one-level dissemination.

3.2.1 Level-Free Advertisement

Level-free dissemination is the most intensive method for vehicular advertising, which allows any nodes to reuse AD_S. Also, the advertising process is simpler than the one-level case in that it does not require a voucher. Rather, vehicles, for making their virtual cash, can return only their collected receipts to the virtual cashier. Here company S, for example, has decided to use level-free dissemination.

S's Ad Distribution Point (ADP) contacts with each approaching vehicle u as follows:

$$S \longrightarrow *: \quad C_S, AD_S, \{AD_S|S\}_{K_{VA}^-}.$$

where C_S is S's certificate. Unlike the three-way handshaking process in one-level case, here each vehicle u only needs to check AD_S's legitimacy. Then, with an incentive in mind, u may start advertising AD_S to any neighboring vehicle v in behalf of S. From now on, everything else is the same as the one-level case, except that vehicle v now can simply reuse AD_S from u (see Figure 2), since virtual cashier VC does not check for a voucher of AD_S for level-free advertisement. Any vehicles advertising AD_S can redeem their collected receipts without vouchers at VC. As in the one-level case, VC also examines whether a redeeming node has never been paid for AD_S at any other cashier before, by inquiring via a VA.

3.2.2 n-Level Advertisement

Such level-free dissemination, however, involves a heavy outlay for advertisement due to its too much redundancy. Thus, in some cases, company S probably wants to set a limit on the number of propagation levels, as a compromise between one-level and level-free case. We refer to this model as n-level advertisement where the only nodes within level-n can reuse the ads. For instance, if $n = 3$ in Figure 2, node x can reuse the ad, but y cannot. Here S, for example, has decided to use n-level dissemination.

We first bring up with a simple design using one-way hash functions, then describe a more refined design which secures n-level dissemination.

Leveraging One-way Hash Chain S's ADP contacts with each approaching vehicle u as follows:

$$S \longrightarrow *: \quad C_S, AD_S, n, \alpha, \{AD_S|S|H^n(\alpha)\}_{K_{VA}^-}.$$

where C_S is S's certificate, n is the number of levels that S sets, α is a random number by S, and $H^n(\alpha)$ is a value applied a one-way hash function H n times on α.[6] Vehicle u checks AD_S's legitimacy by applying H n times on α.

Then, with an incentive in mind, u may advertise AD_S to any neighboring vehicle v, with reducing n by 1 and applying H on α

[6]To prevent DoS attacks, a VA can predefines the maximum value for n.

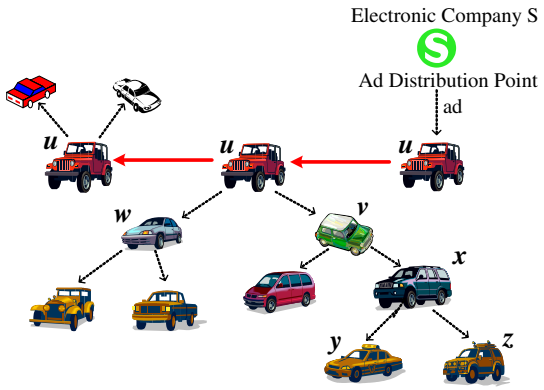

Figure 2: Multi-Level Advertisement.

as follows:

$$u \longrightarrow *: \quad C_u, AD_S, S, n-1, H(\alpha), \{\{AD_S|S|H^n(\alpha)\}_{K^-_{VA}}\}_{K^-_u}.$$
$$v \longrightarrow u: \quad C_v, \{AD_S|u\}_{K^-_v}.$$

where C_u and C_v are u's and v's certificates, respectively. v first checks both u's identity and AD_S's legitimacy by applying H $n-1$ times on $H(\alpha)$. If correct, v may respond to this ad, with an incentive in mind, by providing its receipt for u.

If $n-1$ is still non-zero, v can reuse AD_S, also with reducing $n-1$ by 1 and applying H on $H(\alpha)$, to any neighboring vehicle x as follows:

$$v \longrightarrow *: \quad C_v, AD_S, S, n-2, H^2(\alpha), \{\{AD_S|S|H^n(\alpha)\}_{K^-_{VA}}\}_{K^-_v}.$$
$$x \longrightarrow v: \quad C_x, \{AD_S|v\}_{K^-_x}.$$

where C_v and C_x are v's and x's certificates, respectively.

If $n-2$ is still non-zero, x can reuse AD_S and later submit it to any virtual cashier VC as follows:

$$x \rightarrow VC: \quad C_x, AD_S, S, n-2, H^2(\alpha), \{\{AD_S|S|H^n(\alpha)\}_{K^-_{VA}}\}_{K^-_x}, R_x.$$
$$R_x = (C_y, \{AD_S|x\}_{K^-_y}), (C_z, \{AD_S|x\}_{K^-_z}), ...$$

where R_x is the receipts of AD_S collected by x. VC first checks that $n-2$ is non-zero, and if yes, checks its legitimacy by applying H $n-2$ times on $H^2(\alpha)$.

In this way, AD_S can be reused until it reaches over n levels. This design, however, has some weaknesses. First, it does not have any coercive measures for advertising nodes to reduce their permissible levels by 1 with applying H; these nodes just *voluntarily* follow it for the system in this design. If most nodes reuse AD_S without reducing the levels, it works almost like the level-free advertisement. Second, malicious user w, for example, can throw any permissible value, which does not reach $H^n(\alpha)$ yet (e.g. α,..., $H^{n-1}(\alpha)$), open to the public so that any vehicles even beyond level-n reuse AD_S without directly contacting with w. This also leads to violation of n-level dissemination.

Onion Voucher As an alternative to a one-way hash chain based design, S can manage n-level dissemination by using *onion voucher*. We name it "onion voucher", inspired by Onion routing [37] which disguises the traffic flow, but they are different in goals as well as designs. Figure 3 shows an example of onion vouchers.

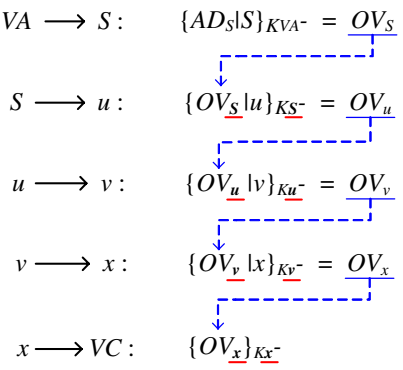

Figure 3: An example of Onion Vouchers.

S's ADP contacts with each approaching vehicle u as follows:

$$S \longrightarrow *: \quad C_S, AD_S, \{AD_S|S\}_{K^-_{VA}}.$$
$$u \longrightarrow S: \quad C_u, \{AD_S\}_{K^-_u}.$$
$$S \longrightarrow u: \quad \{\{AD_S|S\}_{K^-_{VA}}|u\}_{K^-_S}.$$

where C_S and C_u are S's and u's certificates, respectively. Here, after verifying each other, S provides u with an onion voucher $OV_u = \{\{AD_S|S\}_{K^-_{VA}}|u\}_{K^-_S}$[7] for u's exclusive use.

Then, with an incentive in mind, u may advertise AD_S to any neighboring vehicle v as follows:

$$u \longrightarrow *: \quad C_u, AD_S, S, \{\{AD_S|S\}_{K^-_{VA}}\}_{K^-_u}.$$
$$v \longrightarrow u: \quad C_v, \{AD_S|u\}_{K^-_v}.$$
$$u \longrightarrow v: \quad C_S, \{OV_u|v\}_{K^-_u}.$$
$$OV_u = \{\{AD_S|S\}_{K^-_{VA}}|u\}_{K^-_S}$$

where OV_u is u's onion voucher, C_S, C_u, and C_v are S's, u's, and v's certificates, respectively. Exchange of the first two messages is the same as the one-level case. Then if n specified in AD_S is bigger than the number of nodes (except S) included in OV_u, u now provides v's onion voucher $OV_v = \{OV_u|v\}_{K^-_u}$ along with certificates of the nodes (except u and v) included in OV_v. Then v verifies it, by regenerating OV_v with the information provided by u.

Note that v's onion voucher $OV_v = \{OV_u|v\}_{K^-_u}$, in this example, can be properly generated exclusively by u. Even if OV_u is open to the public, any other nodes except u fail to generate new legitimate onion vouchers based on OV_u without u's digital signature on it (i.e. $OV_v = \{OV_u|v\}_{K^-_u}$), which keeps consistent linkage inside an onion voucher.

If OV_v from u is legitimate, v can reuse AD_S to any neighboring vehicle x as follows:

$$v \longrightarrow *: \quad C_v, AD_S, S, \{\{AD_S|S\}_{K^-_{VA}}\}_{K^-_v}.$$
$$x \longrightarrow v: \quad C_x, \{AD_S|v\}_{K^-_x}.$$
$$v \longrightarrow x: \quad C_S, C_u, \{OV_v|x\}_{K^-_v}.$$
$$OV_v = \{\{\{AD_S|S\}_{K^-_{VA}}|u\}_{K^-_S}|v\}_{K^-_u}$$

where OV_v is v's onion voucher, C_S, C_u, C_v, and C_x are S's, u's, v's, and x's certificates, respectively. We note that, if n is

[7]Later, u redeems its digital signature on this onion voucher $\{OV_u\}_{K^-_u}$, along with its collected receipts.

smaller (or equal) than the number of nodes in OV_v, v does not need to provide x's onion voucher. Unlike the three-way handshake initiated by the nodes within level $n-1$, the nodes that belong to level n perform the two-way message exchange advertisement.

If OV_x (i.e. $\{OV_v|x\}_{K_v^-}$) provided by v is still legitimate, x can reuse AD_S and later submit the collected receipts along with its onion voucher OV_x to any virtual cashier VC as follows:

$$\begin{aligned} x \rightarrow VC: \quad & C_x, AD_S, S, C_{OV_x}, \{OV_x\}_{K_x^-}, R_x. \\ C_{OV_x} &= C_S, C_u, C_v \\ OV_x &= \{\{\{\{AD_S|S\}_{K_{VA}^-}|u\}_{K_S^-}|v\}_{K_u^-}|x\}_{K_v^-} \\ R_x &= (C_y, \{AD_S|x\}_{K_y^-}), (C_z, \{AD_S|x\}_{K_z^-}), ... \end{aligned}$$

where C_{OV_x} is certificates of the nodes (except x) included in OV_x, OV_x is x's onion voucher, R_x is the receipts of AD_S collected by x. VC checks whether the number of nodes (except S) included in OV_x is not bigger than n, by regenerating OV_x with the information provided by x. Without a legitimate onion voucher, x fails to get its virtual cash for AD_S at any virtual cashier.

Using onion vouchers, S is able to secure n-level dissemination, in that malicious user w, who tries to throw its voucher OV_w open to the public with the intention of disrupting the n-level advertisement, now has to directly contact with each node z in the network (i.e. w provides z's onion voucher, $OV_z = \{OV_w|z\}_{K_w^-}$); this does not violate n-level dissemination, rather such behavior is desirable for advertisement. However, unlike a one-way hash based one, this design requires three-way handshake message exchange between an advertising vehicle and a receipt-providing one. Also, there is no coercive measure for advertising nodes to generate/provide authentic onion vouchers for receipt-providing nodes in the design. We will discuss this issue in Section 4.5.

4. EVALUATIONS

In this section, we evaluate the performance of ad dissemination designs in SSD in terms of communication, storage, and computational cost, and analyze the incentives and the security of SSD. We then present the simulation results of SSD.

4.1 Communication Cost

SSD involves both inter-vehicle and vehicle-roadside communications. Compared with the highly frequent car-to-car interaction for ad dissemination, communication with roadside infrastructure, however, is not common in SSD, in that vehicles contact only with ADP in the beginning and VC in the end. Hence, we focus on evaluating inter-vehicle communication cost.

In one-level and level-free dissemination models, advertising nodes use the same message format which includes sender's certificate, ad content, ad provider ID, and sender's signature on ad-permit. Although a typical X.509 v3 certificate [16] within IETF is about 1 Kbyte, reduced-size certificates (e.g. a streamlined certificate format [15]) occupy less than 200 bytes. Furthermore, [34] shows that, for public key and signature size, Elliptic Curve Cryptography (ECC) and NTRUSign are the most acceptable PKI implementations for VANETs; ECC (28 bytes) is more compact but slower than NTRU (197 bytes), as in Table 1[8]. Assuming to utilize ECC and 84-byte ECDSA (Elliptic Curve Digital Signature Algorithm)-signed certificate, for x bytes of ad content and the byte size l of ad provider ID, the total message size is $(112 + l + x)$ bytes.

One-way hash chain based n-level model adds two more fields in the above format: the value indicating permissible level (1 byte) and its corresponding hash value (20 bytes in SHA-1). The resulting message size is $(133 + l + x)$ bytes. Thus, in all above cases, considering that l is a small value (e.g. 4 or 8 bytes), the total size of an advertising message mainly depends on x, the size of an ad content itself. We will discuss the upper bound of x, according to the ECC execution time in Section 4.3.

[8]We extract the figures in the table from [34].

Table 1: A comparison between ECC and NTRU on a Pentium II 400MHz workstation.

PKCS	Key size	Signing	Verification
ECC	28 bytes	3.255 ms	7.617 ms
NTRU	197 bytes	1.587 ms	1.488 ms

On the other hand, onion voucher based n-level model performs the three-way handshake advertising process, which incurs the transmission of two separate messages from an advertising node. The first message format is identical to that of the above models, yet the last message format is dependent on the level to which the advertising node belongs. It contains onion voucher OV and certificates of the nodes (except communicating parties) included in OV. The resulting size of the second message from an advertising node in level d $(d < n)$, for example, is $(d \times 84 + 28)$ bytes.

Therefore, as a node belongs to lower level (i.e. as d increases), the size of the second message increases. However, VA can set the maximum value for n, so that d has the upper bound of the propagation level. All the above cases, an ad-receiving node only provides a message containing its certificate and signed receipt so that the resulting message size is 112 bytes.

4.2 Storage Requirement

In one-level ad model, each advertising vehicle needs to store ad content, ad permit, voucher, and all the collected receipts and their corresponding certificates. Thus, for k collected receipts and x bytes of ad content, the total storage requirement is $(k \times 112 + x + 56)$ bytes. The storage requirements of level-free and one-way hash chain based n-level model are similar to one-level case, except that the former does not require 28-byte voucher, and the latter requires additional 1-byte level-indicator and 20-byte hash value.

In onion voucher based n-level model, an advertising node needs to also store a list of certificates included in an onion voucher so that, for a node in level d $(d < n)$, the resulting storage requirement is $(d \times 84 + k \times 112 + x + 28)$ bytes.

In addition, each vehicle may have multiple kinds of advertisements at a time. For example, if an advertising node has g kinds of one-level ads (each ad_i with x_i bytes of ad content and k_i collected receipts) and h kinds of onion voucher based n-level ads (each ad_i with y_i bytes of ad content, d_i level, and p_i collected receipts), then the overall storage requirement (bytes) at the moment is

$$\sum_{i=1}^{g}(k_i \times 112 + x_i + 56) + \sum_{i=1}^{h}(d_i \times 84 + p_i \times 112 + y_i + 28).$$

This shows that the more a node collects the receipts, the required memory storage gets larger. Therefore, considering that d_i is much lower than p_i in most cases, the overall storage requirement mainly depends on the number of the collected receipts.

4.3 Computation Overhead

To justify the choice of ECC in our ad models, we express the computation overhead, according to ECC implementation, inspired from [34] that presents numerical upper bounds for NTRU.

All of our ad models (except for onion voucher based model) em-

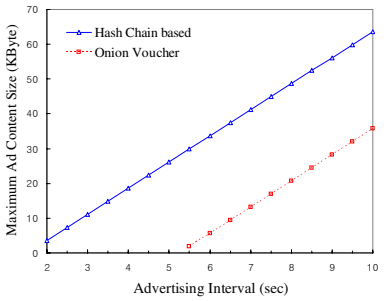

Figure 4: Ad content size (100 neighbors)

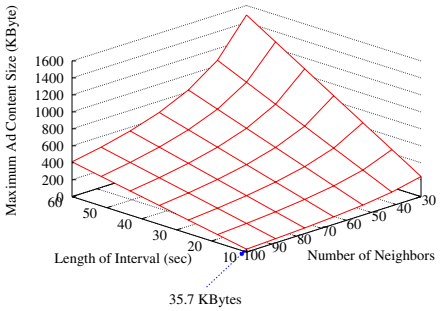

Figure 5: Upper bound of ad content size

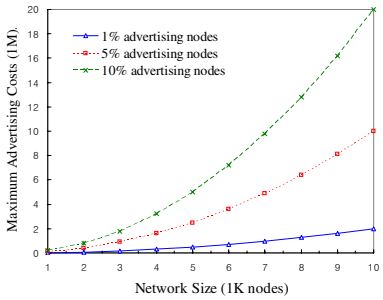

Figure 6: Advertising costs ($\gamma = \delta = 1$)

ploy two-message exchange among vehicles; an advertising node broadcasts an ad, then each receiving node sends back its signed receipt. Each vehicle may receive multiple ads from its neighbors and may want to respond back to all of them, while it may also send out its ads during the same period. Such receipt-providing node's operation involves both verification of the received ad and generation of signature on the receipt. Hence, verifying and signing times are both critical in this case.

Consider a scenario where vehicle u has 100 neighbors within its communication range (<250m), and all the neighbors send out their ads at regular intervals of r ms. For the worse-case condition, all the ads here are of one-way hash chain based n-level model with 8-byte ad-provider ID and x bytes of the ad content size: $(141+x)$ bytes/msg. Thus, assuming that every node sends out its ads and responds back to all others' ads, the resulting system throughput here is $\frac{1000}{r}$ msg/sec $\times$ {[100 cars $\times$ $(141+x)$ bytes/msg] + [100 $\times$ 100 cars $\times$ 112 bytes/msg]}

$$= 8 \cdot \{(141+x) + (112 \times 10^2)\}/(r \cdot 10) \text{ Mbps.} \quad (1)$$

Since node u needs to process all of such incoming ads while sending out its ads at intervals of r ms, the permissible processing time per each incoming ad should be larger than ECC's verifying[9] plus signing delay specified in Table 1. Hence, r ms/100 > 18.45 ms, so $r > 1.845$ sec. By setting 6 Mbps (minimum data rate in DSRC) to Equation (1), we now can derive the upper bound of ad content size x with respect to interval r. Figure 4 shows that the upper bound of x, according to r in this 100-neighbor scenario.

On the other hand, in the case of onion voucher based model, due to three-way handshake process among vehicles, the resulting system throughput of the above scenario is $\frac{1000}{r}$ msg/sec $\times$ {[100 cars $\times$ $(120+x)$ bytes/msg] + [100 $\times$ 100 cars $\times$ 112 bytes/msg] + [100 $\times$ 100 cars $\times$ 280 bytes/msg]}

$$= 8 \cdot \{(120+x) + (392 \times 10^2)\}/(r \cdot 10) \text{ Mbps.} \quad (2)$$

Here we assume that every node here belongs to level 3 so that its second message size is 280 bytes. Vehicle u needs to process all of such incoming ads as well as all the incoming receipts to generate onion vouchers, while sending out its ads at intervals of r ms. Such voucher-generating operation involves verification of the receipt and generation of signature on a new onion voucher. Hence, the permissible processing time per each neighbor's ad and receipt should be larger than ECC's ad-processing time (18.45 ms) plus receipt-processing time (10.87 ms) drawn from figures in Table 1.

[9]Ad message actually incurs two rounds of verification process, due to a doubly-signed ad-permit.

Therefore, r ms/100 > 29.32 ms, so $r > 2.932$ sec. With 6 Mbps to Equation (2), the upper bound of ad content size x can be obtained. As seen in Figure 4, to meet the acceptable size of ad content for this 100-neighbor scenario, the feasible advertising interval should be at least 5.5 sec.

Among our ad models, the onion voucher ad model turns out to be the worse case outcome, on mainly account of three-way handshake among vehicles. Thus, we can consider the onion voucher model's ad content size as the ultimate upper bound of all our ad models'. Figure 5 shows the ultimate upper bound of ad content size, with varying the length of interval and the number of neighbors within communication range. We note that a node with h different kinds of ads sends out each ad every r ms (e.g. round-robin), so that the average advertising interval between the same kind of ad is $h \times r$ ms.

4.4 The Incentive Perspective

Different from the incentive schemes [1, 5, 13, 20, 42, 48] for multi-hop packet delivery, our vehicular ad system has a single-hop broadcast characteristic. SSD gives an incentive to both advertising nodes and receipt-providing nodes, by means of virtual cash. Each individual node u benefits from advertising ads, which increases the chance of obtaining more receipts from ad-receiving nodes. u also benefits from providing receipts for any neighboring advertising node v, since once those receipts are successfully redeemed by v at any virtual cashier, u will eventually get its share. Thus, when the rewards (i.e. the amount of virtual cash) per each redeemed receipt are γ and δ for a redeeming node and a receipt-providing node respectively, the total (potential) gain for each node u is (*the number of the collected receipts* $\times$ γ) + (*the number of the receipts generated* $\times$ δ). As for an advertising company's costs, in a network of size N with q advertising nodes, the maximum advertising costs C_{max} (i.e. virtual cash) invoked by a certain ad is:

$$C_{max} = q \cdot \{(N-1) \cdot \gamma + (q-1) \cdot \delta\} + (N-q)(q-1) \cdot \delta$$
$$\approx q \cdot N \cdot (\gamma + \delta)$$

Figure 6 shows C_{max} when $\gamma = \delta = 1$. As expected, the upper bound of advertising costs increases with the number of ad-forwarding nodes as well as the network size.

We note that the actual gain for each node u, however, can be somewhat lower than the above potential gain. This is because (i) u can hardly encounter every node in the network, rather it may obtain receipts from a subset of N (ii) The virtual cashier also rejects, if any, invalid receipts among u's collected ones (iii) Some nodes for which u has provided receipts may fail to redeem their collected receipts within the specified due date.

There are two possible way for setting the values of γ and δ; a VA sets the fixed values for all the vehicular ads, or advertising companies set their own values. In the former case, all ads are treated equally, but in the latter some ads may have priorities over others. One side-effect of such priorities is the *starvation* of non-high priority ads, since most users may want to keep advertising high-valued ads.

In summary, the incentive property of SSD is that the rewards are always given to both redeeming node u and each receipt-providing node v. If u fails to get its virtual cash, v cannot get its share for the receipt, and vice versa.

4.5 The Security of Signature-Seeking Drive

SSD makes use of a PKI for secure incentives in vehicular ad system. Utilizing a PKI relies on the assistance of the centralized administration (e.g. CA, VA, etc), which conflicts the self-organized and distributed nature of mobile ad-hoc networks. However, to implement primary applications (e.g. road safety, traffic management, payment services, etc) of VANETs, a PKI with the centrally-managed infrastructure is an essential part of the networks.

Malicious Users in Collusion SSD can prevent the colluding nodes from sharing/fabricating their vouchers or collected receipts, since each of those is cryptographically tied to the only one holder's identity. Also, malicious nodes fail to launch DoS attacks (e.g. generating/disseminating dummy ads), since receiving nodes discard such unauthorized ads without a VA's signature on it. The integrity of ad content is protected so that they cannot modify the ad (e.g. free-riding attack [1]). However, once malicious users cheat a VA into issuing ad-permits with fake identities, they are able to disrupt the system without being charged for what they invoke. Thus, the management of a VA as well a CA is the critical issue in our model, yet still a challenging problem in reality.

Securing n-level Advertisement Malicious user w can spoil the one-way hash chain based n-level model by exposing any permissible value open to the public, so that any node z is able to circumvent the n-level restriction. Onion voucher can thwart such attacks. To get a legitimate onion voucher ($OV_z = \{OV_w|z\}_{K_w^-}$) with w's signature on it, z has to directly contact with w. However, if w throws its private key K_w^- along with certificate C_w and voucher OV_w open to the public, any node z is now able to generate a valid voucher by itself, which infringes the n-level model.

Unfortunately, we do not provide a proactive countermeasure against this type of *disclosure* attacks. However, we point out that such disclosure ironically has to be "not open to the public", since the exposure of a private key with certificate reveals the key owner's identity. If such disclosure is detected, the authorities can identify the attacker and may notify the virtual cashier for rejecting all the vouchers associated with it. Each onion voucher keeps *propagation-level history* (e.g. in Figure 3, OV_x: $S \rightarrow u \rightarrow v \rightarrow x$) in it, and such traceable property facilitates the filtering of attacker-involved vouchers.

Another problem of the onion voucher design is that, once obtaining receipts from ad-receiving node x, advertising node v may refuse to generate/provide an onion voucher for x. As an active solution for this problem, we can give an extra credit for nodes providing onion vouchers. For example, when redeeming OV_x in Figure 3, the last node (i.e. v) in OV_x likewise gets its share. Thus, v can benefit from providing OV_x, yet x redeems only one voucher among multiple ones so that the expected gain of onion voucher providing is relatively lower than that of receipt providing.

Anonymous Keys & Receipt Verification Preserving drivers' privacy is the practical issue in VANETs. Specifically in our SSD, each advertising node collects receiving nodes' certificates as well as their signed receipts, which raises privacy concerns. Thus, to protect privacy against unauthorized observers, each vehicle is likely to have its temporary IDs (i.e. a set of anonymous public/private key pairs), whose certificates have short lifetimes (e.g. a few minutes with the reanonymizer [32], or several days with the preloading scheme [34]), so that each vehicle changes its signing key periodically. This helps prevent third parties from tracking the real identities of communicating vehicles.

However, in SSD, advertising node u may obtain the distinct (signed with different keys) receipts (e.g. $\{AD_I|u\}_{K_{v1}^-}$, $\{AD_I|u\}_{K_{v2}^-}$, etc) from node v, if encountered multiple times. Such multiple receipts from v should be counted as "a single receipt" at virtual cashier VC. Also, advertising node u likewise uses anonymous keys. Thus, when u redeems its collected receipts R_u, the receipt-verification process at VC involves the following: (i) Examining that each signing key in R_u was valid at the time of being used. Since unable to track the actual signing time, VC checks whether the valid period of each key in R_u is within the ad's term of validity, as a minimum requirement. (ii) Counting valid receipts from the actually distinct nodes (i.e. one receipt from one node), which requires mapping of anonymous public keys into the actual identities of the vehicles. To avoid abuse, such ID matching capability should be shared among multiple authorities [34], for example, VC and a VA can share the secret for retrieving the actual IDs. Another possible approach is the layered auditing [8], where the front-end (e.g. VC) and the back-end (e.g. VA) authorities have different privileges but collaborate for ID retrieval.

Miscellaneous Issues When the dissemination of ads within a certain area is strictly required, independent from setting the dissemination level, we can rely on secure positioning schemes [6, 7] so that the only vehicles within the specified area can generate valid receipts with certified location information in them. VC rejects the receipts without such information, and advertising nodes cannot get a valid receipt outside the target area.

Malicious node w may launch DoS attacks by incessantly sending out its ads without any interval, on purpose to provoke *receipt-implosion* toward w itself. To alleviate these attacks, each node keeps a cache [14, 19] for the recently-received ads (along with the senders' IDs), so that it can suppress the same copies of previously-generated receipts. Since the above temporary certificates' lifetimes are longer than the desired advertising interval, w cannot change its pseudo IDs within the interval. Still, w can launch other types of DoS attacks such as physical-layer jamming attacks [10, 29, 43], which are beyond the scope of this paper.

Compared with an acceptable latency and a high data rate of DSRC, its message delivery ratio turns out to be still lacking [45]. Especially in high speed traffic situation where vehicles moving in opposite directions contact for an instant, the three-way handshake in the onion voucher model may incur a great deal of message loss. However, we note that, unlike safety-related applications in VANETs, ad content itself is not critical information, nor of real-time constraints, rather a soft state[10].

4.6 Simulations

To further evaluate the performance of our ad models, we run simulations of SSD on a network simulator, *ns-2* [28].

4.6.1 Simulation Model

In our simulation, we use the mobility model [38] specifically designed for VANETs, where the real city map data is used as an input

[10]Soft state [9] means the state information could be lost without permanent disruption of the service features being used.

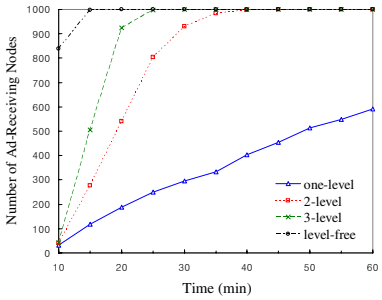

Figure 7 Number of ad-receiving nodes with 1% Level-1 nodes

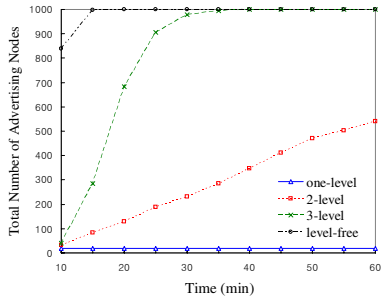

Figure 8 Number of advertising nodes with 1% Level-1 nodes

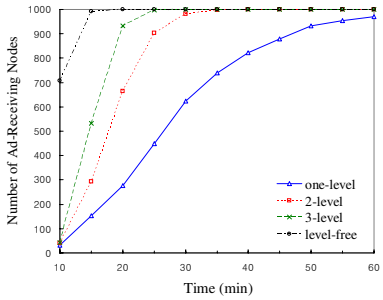

Figure 9 Number of ad-receiving nodes with 5% Level-1 nodes

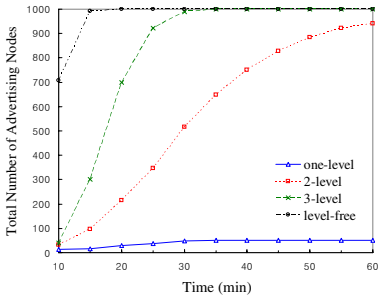

Figure 10 Number of advertising nodes with 5% Level-1 nodes

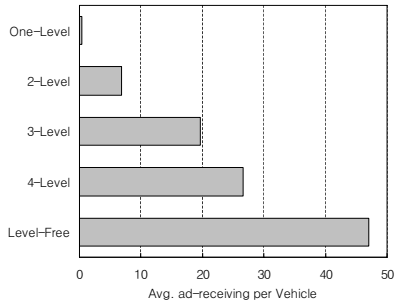

Figure 11 Average number of duplicate Ads received per Vehicle

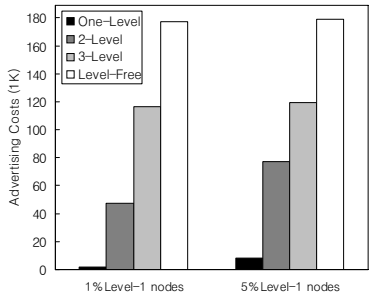

Figure 12 Advertising costs ($\gamma = \delta = 1$)

for ns-2 simulator. We extract/convert a street map of $4 \times 4Km^2$ Westwood area around UCLA campus, from the US Census Bureau's TIGER (Topologically Integrated Geographic Encoding and Referencing) database [39], into our simulation environment. The resulting target area contains 1-hour movement pattern of 1000 vehicles that travel inside the area. We position an ADP in the center of the area. It distributes an ad to the predefined number of passing-by vehicles, which are level-1 nodes in the dissemination hierarchy. Simulation experiments are conducted with 1% and 5% of the total population as level-1 nodes, varying the ad model from one-level to level-free advertisement. In the simulations, each advertising node sends out its ad every 30 sec with $250m$ transmission range.

It is worth noting that there are some unrealistic aspects of the mobility model in the simulation as follows: (i) Each node selects the random starting/destination points in the target area. In real-world scenario, instead of random movement, each individual or a group of cars are likely to have certain traffic patterns (e.g. commuting). (ii) All the nodes in the network are always moving within the area. However, in reality, the majority of private cars show the occasional movement patterns (e.g. parked for several hours), and moreover many cars are coming into or leaving the area. (iii) They have no traffic control so that every vehicle in the area moves always at the rate of a street speed limit.

In addition, the number of cars in our simulation model is much lower than that of the real world. In fact, there are more than 10,000 cars in Westwood area, yet none of existing event-driven ad hoc network simulators, to the best of our knowledge, is capable of processing more than tens of thousands of nodes.

4.6.2 Simulation Results

We first measure the ad-coverage (i.e. the number of ad-receiving nodes in the area) and the total number of advertising nodes with one-level, onion voucher based n-level, and level-free advertisement.

Figure 7 and 9 illustrate the simulation results of the ad-coverage with 1% and 5% of level-1 nodes, respectively. As expected, as the advertisement level n increases, an ad tends to propagate faster over the network. Especially, the level-free advertisement shows the most rapid ad-coverage rate (100% ad-coverage within 15 min). However, in real-world scenario, the rate of such ad-coverage slows down for the reasons described earlier (e.g. at any moment, a large percentage of cars in the area are inactive for being parked).

Figure 8 and 10 show the total number of advertising nodes in the same experiments. Unlike the one-level advertisement whose number of advertising nodes remains at 10 (1%) and 50 (5%) respectively, the others have an increasing number of advertising nodes over time. Since the resulting number of advertising nodes in the network is strongly related with an advertising company's costs (as in Figure 6), the level-free ad model which displays the sharpest slope in Figure 8 and 10 incurs the heaviest outlay for advertisement. This is also demonstrated in Figure 11, which plots the average number of the duplicate ads received per vehicle within 30 min. Such reception of the duplicate ad is desirable for advertisement, since the repeated delivery of ads to the potential customers enhances the effect of the commercial ads. We can infer that, from the simulation results in Figure 7 $\sim$ 11, as n grows in real-world scenario, each ad model (except the one-level case) entails the higher increasing rate of the number of advertising nodes, which results in higher advertising costs.

Figure 12 provides a more detailed view on each ad model's advertising costs, assuming that this ad is valid for 1 hour and every advertising node successfully redeems its receipts. As can be seen from the figure, level-free model shows the highest advertising costs followed by n-level models. We can also find that each

ad model's advertising costs are proportional to its average number of duplicate ads received per a vehicle (in Figure 11), in that the number of the receipts generated at each node increases with the number of the ads received.

In summary, through simulation experiments, the level-free model outperforms the others with respect to the ad-coverage rate and the average number of the duplicate ads received, yet it shows the highest increasing rate of the number of advertising nodes. We can see that n-level model is a gradual compromise between one-level and level-free advertisement.

5. CONCLUSIONS AND FUTURE WORK

In this paper we have proposed a potential and promising application scenario, the dissemination of commercial advertisements in VANETs. With both selfish users (incentives) and malicious users (security) into account, we have presented Signature-Seeking Drive (SSD), a secure incentive framework for commercial ad dissemination in VANETs. Unlike currently proposed incentive systems, SSD does not rely on tamper-proof hardware or game theoretic approaches, but leverages a PKI to provide secure incentives for cooperative nodes. With a set of ad dissemination designs proposed, we have demonstrated that our SSD is robust in both incentive and security perspectives.

In future work, we plan on further extending our framework from the domain of vehicles to that of pedestrians with hand-held computers (e.g. PDAs, Smartphones, etc). We will also seek to develop an incentive system in which an ad-provider can give the rewards directly to cooperative users without the intervention of a VA. Finally, we would like to perform an extensive experimental evaluation of SSD within a realistic experimental environment.

6. REFERENCES

[1] N. Ben Salem, L. Buttyan, J.-P. Hubaux and M. Jakobsson. A Charging and Rewarding Scheme for Packet Forwarding. *ACM MobiHoc*, 2003.

[2] S. Buchegger and J.-Y. L. Boudec. Nodes Bearing Grudges: Towards Routing Security, Fairness, and Robustness in Mobile Ad hoc Networks. *Euromicro Workshop on Parallel, Distributed and Network-based Processing*, 2002.

[3] S. Buchegger and J.-Y. L. Boudec. Performance Analysis of the CONFIDANT Protocol: Cooperation of Nodes - Fairness in Dynamic Ad-hoc Networks. *ACM MobiHoc*, 2002.

[4] L. Buttyan and J. P. Hubaux. Enforcing Service Availability in Mobile Ad-hoc WANs. *ACM MobiHoc*, 2000.

[5] L. Buttyan and J. P. Hubaux. Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks. *ACM Journal for Mobile Networks (MONET)*, 2002.

[6] S. Čapkun, L. Buttyan, and J.-P. Hubaux. SECTOR: Secure Tracking of Node Encounters in Multi-Hop Wireless Networks. *ACM Workshop SASN*, 2003.

[7] S. Čapkun and J.-P. Hubaux. Secure Positioning of Wireless Devices with Application to Sensor Networks. *IEEE INFOCOM*, 2005.

[8] J. Y. Choi, M. Jakobsson, and S. Wetzel. Balancing Auditability and Privacy in Vehicular Networks. *ACM Workshop Q2SWinet*, 2005.

[9] D. D. Clark. The Design Philosophy of the DARPA Internet Protocols. *ACM SIGCOMM Computer Communication Review*, vol. 25 issue 1, 1995.

[10] J. Deng, R. Han, and S. Mishra. Defending against Path-based DoS Attacks in Wireless Sensor Networks. *ACM Workshop SASN*, 2005.

[11] 5.9 GHz Dedicated Short Range Communications (DSRC). `http://grouper.ieee.org/groups/scc32/dsrc/index.html`.

[12] D. Eastlake and P. Jones. US Secure Hash Algorithm 1 (SHA1). *RFC 3174*, September 2001.

[13] M. Felegyhazi, J.-P. Hubaux, and L. Buttyan. Nash Equilibria of Packet Forwarding Strategies in Wireless Ad Hoc Networks. *IEEE Transactions on Mobile Computing (TMC)*, 2006.

[14] W.R. Heinzelman, J. Kulik, and H. Balakrishnan. Adaptive Protocols for Information Dissemination in Wireless Sensor Networks. *ACM MobiCom*, 1999.

[15] G. Horn and B. Preneel. Authentication and Payment in Future Mobile Systems., *Journal of Computer Security*, vol.8, pp.183-207, 2000.

[16] R. Housley, W. Ford, W. Polk, and D. Solo. Internet X.509 Public Key Infrastructure: Certificate and CRL Profile. *RFC 3280*, April 2002.

[17] E. Huang, J. Crowcroft, and I. Wassell. Rethinking Incentives for Mobile Ad Hoc Networks. *ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, 2004.

[18] J.-P. Hubaux, S. Capkun, and J. Luo. The Security and Privacy of Smart Vehicles. *IEEE Security & Privacy Magazine*, pp. 49–55, 2004.

[19] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. *ACM MobiCom*, 2000.

[20] M. Jakobsson, J.P. Hubaux, and L. Buttyan. A Micro-Payment Scheme Encouraging Collaboration in Multi-Hop Cellular Networks. *Financial Crypto*, 2003.

[21] I. Khalil, S. Bagchi, and N. B. Shroff. LITEWORP: A Lightweight Countermeasure for the Wormhole Attack in Multihop Wireless Networks. *DSN*, 2005.

[22] J. Kangasharju and A. Heinemann. Incentives for Electronic Coupon Systems. *ACM Workshop MobiShare*, 2006.

[23] J.-Y. Lee and R.A. Scholtz. Ranging in a Dense Multipath Environment Using a UWB Radio Link. *IEEE Journal of Selected Areas in Communications*, 2002.

[24] U. Lee, J.-S. Park, E. Amir, and M. Gerla. FleaNet: A Virtual Market Place on Vehicular Networks. *V2VCOM*, 2006.

[25] Y. Liu and Y. R. Yang. Reputation Propagation and Agreement in Mobile Ad-hoc Networks. *IEEE WCNC*, 2003.

[26] J. Luo and J.-P. Hubaux. A Survey of Research in Inter-Vehicle Communications. *Securing Current and Future Automotive IT Applications*, pp 111–122, Springer-Verlag, 2005.

[27] S. Marti, T.J. Giuli, K. Lai, and M. Baker. Mitigating Routing Misbehavior in Mobile Ad Hoc Networks. *ACM MobiCom*, 2000.

[28] S. McCanne and S. Floyd. ns Network Simulator. `http://www.isi.edu/nsnam/ns/`

[29] J.M. McCune, E. Shi, A. Perrig, and M.K. Reiter. Detection of Denial-of-Message Attacks on Sensor Network Broadcasts. *IEEE Security & Privacy*, 2005.

[30] A. Nandan, S. Das, B. Zhou, G. Pau, and M. Gerla. AdTorrent: Digital Billboards for Vehicular Networks. *V2VCOM*, 2005.

[31] V. Naumov, R. Baumann, and T Gross. An Evaluation of Inter-Vehicle Ad Hoc Networks Based on Realistic Vehicular Traces. *ACM MobiHoc*, 2006.

[32] B. Parno and A. Perrig. Challenges in Securing Vehicular Networks. *HotNets-IV*, 2005.

[33] C. Passmann, C. Brenzel, and R. Meschenmoser. Wireless Vehicle to Vehicle Warning System. *SAE 2000 World Congress*, 2002.

[34] M. Raya and J.-P. Hubaux. The security of Vehicular Ad Hoc Networks. *ACM Workshop SASN*, 2005.

[35] M. Raya, D. Jungels, P. Papadimitratos, I. Aad, and J.-P. Hubaux. Certificate Revocation in Vehicular Networks. *Technical Report LCA-Report*-2006-006, 2006.

[36] D. Reichardt, M. Miglietta, L. Moretti, P. Morsink, and W. Schulz. CarTALK 2000 - Safe and Comfortable Driving Based upon Inter-Vehicle-Communication. *IEEE Intelligent Vehicle Symposium*, 2002.

[37] M. Reed, P. Syverson, and D. Goldschlag. Anonymous Connections and Onion Routing. *IEEE Journal on Selected Areas in Communication Special Issue on Copyright and Privacy Protection*, 1998.

[38] A. K. Saha and D. B. Johnson. Modeling Mobility for Vehicular Ad Hoc Networks. *ACM Workshop VANET*, 2004.

[39] U.S. Census Bureau. 2005 Second Edition TIGER/Line Files. `http://www.census.gov/geo/www/tiger/tiger2005se/tgr2005se.html`

[40] H. Vogt, F. C. Gartner, and H. Pagnia. Supporting Fair Exchange in Mobile Environments. *ACM Mobile Networks Journal (MONET)*, 2003.

[41] Q. Xu, T. Mark, J. Ko, and R. Sengupta. Vehicle-to-Vehicle Safety Messaging in DSRC. *ACM Workshop VANET*, 2004.

[42] H.-Y. Wei and R. D. Gitlin. Incentive Mechanism Design for Selfish Hybrid Wireless Relay Networks. *ACM Mobile Networks Journal (MONET)*, 2005.

[43] A. Wood and J. Stankovic. Denial of Service in Sensor Networks. *IEEE Computer*, vol.35, 54-62, Oct. 2002.

[44] X. Yang, J. Liu, F. Zhao, and N. Vaidya. A Vehicle-to-Vehicle Communication Protocol for Cooperative Collision Warning. *MobiQuitous*, 2004.

[45] J. Yin, T. ElBatt, G. Yeung, B. Ryu, S. Habermas, H. Krishnan, and T. Talty. Performance Evaluation of Safety Applications over DSRC Vehicular Ad hoc Networks. *ACM Workshop VANET*, 2004.

[46] A. Zanella, E. Fasolo, C. F. Chiasserini, M. Meo, M. Franceschinis, and M. A. Spirito. Inter-Vehicular Communication Networks: a Survey. *2nd Internal NEWCOM Workshop*, 2006.

[47] M. E. Zarki, S. Mehrotra, G. Tsudik, and N. Venkatasubramanian. Security Issues in a Future Vehicular Network. *EuroWireless*, 2002.

[48] S. Zhong, J. Chen, and Y. R. Yang. Sprite: A Simple, Cheat-Proof, Credit-Based System for Mobile Ad-Hoc Networks. *IEEE INFOCOM*, 2003.