A Survey of Vehicle Attacks and Defenses in Vehicular Ad-Hoc Networks

A Senior Project

presented to

the Faculty of the Computer Science

California Polytechnic State University, San Luis Obispo

In Partial Fulfillment

of the Requirements for the Degree

Enter Your Degree Name; e.g. Bachelor of Science

by

Pierson Yieh

June, 2018

Introduction

With recent technological advances in intelligent transportation systems (ITS), Vehicular Ad-Hoc Networks (VANETs) have seen expanded applications. VANETs are systems of vehicles communicating with each other and roadside infrastructures. One application for VANETs is Vehicle Safety Communication (VSC), which aim to enhance vehicle safety and the driving experience. VSC can be further broken down into vehicle-to-vehicle (V2V) communications and vehicle-to-infrastructure (V2I) communications. An example of V2V communication is platooning, where vehicles closely follow other vehicles aided by wirelessly exchanging steering and acceleration information. Examples of V2I communication are automatic tollbooth payments, where vehicles can automatically send their occupants' payment information to tollbooths without having to stop, and emergency roadside warnings, where roadside infrastructures can broadcast to vehicles information about upcoming dangerous weather and road conditions [1]. Because VANETs use wireless technology for communication, attackers can easily pick up signal packets with proper equipment [2, 3]. This brings into question whether a malicious at-tacker can track users for an extended period of time using these messages, in what we refer to as a tracking attack.

One means of mitigating a tracking attack is using pseudonyms in VANET messages [4]. Pseudonyms are temporary unique identifiers used by vehicles when sending VSC messages that are switched regularly to prevent tracking over long periods of time. Pseudonyms are distributed to vehicles by a trusted third-party Certificate Authority (CA), with vehicles' owners being associated with their vehicles' pseudonyms, for liability reasons [5]. Requiring a CA instead of simply allowing vehicles to generate their own pseudonyms ensures authenticity, prevents pseudonym spoofing, and mitigates pseudonym collisions. While a simple strategy for protecting user privacy in VANETS is to frequently switch pseudonyms, this is not an ideal solution because pseudonym changes are expensive [6]. The cost for changing pseudonyms comes from the limited number of pseudonyms a vehicle can store and the expense, or impossibility, of downloading new pseudonyms. To most effectively change pseudonyms, research has been done into pseudonym change strategies, or algorithms that determine when to best change pseudonyms. While pseudonym change strategies aim to maximize the utility of each pseudonym change, they do not guarantee preserving user privacy.

Previous research has investigated the possibility and effectiveness of tracking wireless devices and vehicles using the beacon messages transmitted periodically by the tracking targets. Recent works showed that an attacker can track many mobile devices with accuracy comparable to GPS and pro-vide high-accuracy trajectory information by simply using inexpensive o - the-shelf equipment [2, 3]. Furthermore, it has been suggested that vehicles' tire pressure monitoring systems (TPMS) contain vulnerabilities that would allow an attacker to perform such tracking attacks [7]. This project is a survey of historical privacy security attacks and defenses in VANETs. I begin with a discussion of recent pseudonym change strategies, then explore privacy attacks that have been proposed in previous research.

Pseudonym Change Strategies

Because of the concern of vehicle tracking through the use of V2X protocols, there have been a number of works exploring effective means of defending against such attacks. A key means of defending against vehicle tracking is the use of pseudonyms [8]. A pseudonym is a unique identifier associated with a signal emitted by a vehicle to allow unique identification for a period of time, as some signal usages require identification, but also changed periodically to prevent long-term tracking. The simple strategy for protecting privacy in VANETs by frequently changing pseudonym is not ideal because pseudonym changes are expensive due to limited pseudonym storage capacities, limited download capabilities for new pseudonyms, and increased costs on various applications [8, 5]. Therefore, effective pseudonym change strategies, algorithms that dictate when a vehicle changes pseudonyms, have been developed to maximize the effectiveness of each pseudonym change, while minimizing the number of changes required.

A common strategy is a periodic change strategy [4], in which the signal protocols dictate how often a vehicle broadcasts a message (heart-rate) and changes its pseudonym (change-rate), with all vehicles having the same heart-rate and change-rate if they are using the same protocol. In a periodic change strategy, vehicles change their pseudonyms every N seconds, using a simple counter to determine when the time limit has been reached. This change strategy is simple and does not require cooperation among neigh-boring vehicles to coordinate pseudonym changes, which prevents malicious adversaries to disrupt privacy gains of pseudonym changes by vehicles within the VANET, a topic which has been investigated in graduate student Nicholas Plewtong's thesis paper [9].

Another common strategy is a random change strategy [4], where a random number is generated before each beacon broadcast. If the generated number is below a predetermined threshold, then the vehicle changes pseudonyms. Similar to the periodic change strategy, the random change strategy is a simple strategy and does not require inter-vehicle cooperation to effectively change pseudonyms.

Another pseudonym change strategy is a synchronous change strategy [4]. The synchronous change strategy can be classified as a type of a position change strategy, where vehicles change pseudonyms when a minimum thresh-old of vehicle density within their proximity is met. The synchronous change strategy requires coordination among the vehicles in the VANET to maximize effectiveness of pseudonym changes. Vehicles that change pseudonyms together would prove to be more di cult for an attacker to associate their new pseudonyms to their respective old pseudonyms due to the lack of distinguishing features and increased number of possible association pairs. The synchronous change strategy has vehicles ready to change pseudonyms c_{min} seconds after their last change. Vehicles indicate their readiness to change by setting a change ag within their broadcast messages. A change is triggered when there are k - 1 other vehicles with their change ag set within the transmission range. This allows for a synchronous pseudonym change by all vehicles ready to change pseudonyms within that given area. If the threshold of k - 1 vehicles is not met within c_{max} seconds after the last change, then the vehicle changes pseudonyms anyways. This is to prevent vehicles from perpetually waiting for the number of vehicles threshold to be met in sparse areas or in areas where this strategy has not been adopted [4]. On its own without the necessary modifications, the synchronous change strategy would be susceptible to malicious adversaries posing as compliant vehicles nodes that aim to disrupt the vehicular privacy of

nodes within the VANET [9]. A similar change strategy called the density-based location privacy scheme has also been proposed that uses the same concept of changing pseudonyms when enough other vehicles are within proximity to make old-new pseudonym associations more di cult [10].

Another type of position change algorithm is a similar status algorithm. In a similar status algorithm, vehicles coordinate pseudonym changes by looking for other vehicles with a similar status as their own, with a pseudonym change occurring when a given number of vehicles with a similar status are within proximity of a respective vehicle. This strategy allows for many different things to be considered features of a vehicle's status such as speed, direction, and number of neighbors [11]. However, this strategy requires the broadcasting of additional descriptive information about the vehicle, which can be used by adversaries to further distinguish vehicles from each other [12].

AMOEBA is another means of defending user privacy that takes advantage of the clustering nature VANETs to prevent malicious attackers [13]. While not a pseudonym change strategy, AMOEBA is a privacy scheme that employs the use of pseudonyms, namely a single group pseudonym representative of all vehicles within a given cluster. Groups are formed by vehicles that move with a similar velocity and relative to each other, and have a fully connected network graph within them to allow for communication among their respective members. A leader of the group is elected at random and broadcasts on behalf of the group members, while other members can remain silent. AMOEBA also defines the use of silent periods to prevent linkability between two locatable broadcasts. Random silent periods are used to prevent trackability. An example would be if a vehicle changes its pseudonym from A to A'. Initially entering a network and broadcasting as A, after having changed and waited a random silent period, it begins broadcasting as A'. If another vehicle had changed pseudonyms from B to B' within that silent period, an attacker may be misled to tracking the neighboring vehicle.

Privacy Attacks

A common attack used to quantitatively compare pseudonym and tracking related defenses is trying to match enter events with their respective exit events [14, 15]. Enter and exit events represent when a vehicle enters or exits a mix zone, respectively. A mix zone is an area outside of an attacker's area of observation, where vehicles can become mixed together without an attacker's knowledge.

The premise of this attack is to match corresponding enter and exit events using previously learned data. The attack is broken into two phases, a learn-ing phase and an attack phase. During the learning phase, the attacker records the number of vehicles that travel between two areas of observation and the average time each vehicle took to make that respective trip. The data learned is limited to what can be learned from observing pseudonyms for the pseudonym’s given lifespan, that is, the time before the vehicle changes pseudonyms. This means that due to vehicles changing pseudonyms, the number of vehicles recorded to have traveled between two locations can be greater than the actual number of vehicles. During the attack phase, the attack rst actively attempts an online attack to match each newly observed exit event to a previously observed enter event. This step simply matches events if they broadcast the same pseudonyms and removes the events from the set of events that are used later. After, all unmatched events are used to create a bipartite graph, with the two distinct sets being exit and enter events. Edges are assigned between an exit and an enter event with a weight equal to the number of vehicles that traveled between the two points where the events were observed divided by the average time each vehicle took., with a penalty added the farther the actual trip time was from the average trip time. If during the learning phase, no trip was recorded between the two locations, then a small weight of 0.1 is given. When the bipartite graph is complete, the solution is a matter of solving the linear sum problem and getting a minimal cost perfect match of the graph. The success of the attack is measured by how many pairs of events were correctly matched.

Attacks have also been developed that take advantage of specific message protocol features. In Examining Privacy in Vehicular Ad-Hoc Networks, the authors break down the privacy vulnerabilities of the DSRC protocol stack [12]. Within SAE J2735, the standard message structure for DSRC messages, a vehicle is identified by a 4-byte temporary identifier pseudonym, while also broadcasting information such as the vehicle's GPS coordinates, motion information, and vehicle size that can be seen by anyone. The authors argue that DSRC messages leak enough information that an attacker can circumvent the temporary nature of the pseudonym and track vehicles despite switching pseudonyms, as well as linking pseudonyms to the actual vehicle or owner. The authors claim that attackers can use statistical methods, similar to my attack, to track vehicles regardless of pseudonym switches. The main means of doing so is when a vehicle switches pseudonyms, an attacker will see a pseudonym no longer transmitting, and a new pseudonym begins transmit-ting within close proximity. The effectiveness of this attack increases with greater coverage and can also take advantage of the descriptive information contained in DSRC messages to better associate pseudonyms in the case of multiple vehicles simultaneously switching pseudonyms. I use a similar method in one of my matrix constructions, but I am limited to the pseudonym information of a new pseudonym appearing shortly after an old pseudonym disappears. The

authors then claim that by using this location information, attackers can link pseudonyms back to the vehicle owners. Aside from the direct linkage of pseudonyms to their owners by gaining access to the pseudonym database, an attacker can use the location information to build a pro le for certain vehicles. Knowing where vehicles stop and go at what times, an attacker can correlate points of interest to buildings and times visited to possibly discover where the user works and lives. After narrowing down buildings, an attacker can further pinpoint the user by performing a lookup of owners and occupants of the buildings for residency and employee directories for businesses in that area.

The authors describe a more general version of the previously mentioned synchronous pseudonym strategy [4] as an effective pseudonym change strategy to prevent tracking, although the use of a similar status algorithm[11] would likely be even more effective, as DSRC messages already contain descriptive information about a vehicle. They also cite the use of a group pseudonym as a means of defending privacy [13], though they do note that an attacker can attack the point when a vehicle leaves a cluster and enters a new one to learn information.

An attack developed by the authors of [16] aims to associate a large set of collected anonymous location samples to anonymous location pro les using the established Multiple Hypothesis Tracking (MHT) algorithm to track vehicles' locations over an extended period of time. MHT addresses the data association problem by generating a set of data associations hypotheses every time a new set of measurements arrives, with each hypothesis being a possible association of a measurement with a target. The probability for each hypothesis to be correct is calculated and the highest probability is chosen to be the solution. MHT relies on Kalman filters [17] to estimate the state variables of each target: position and velocity. This attack assumes the role of a passive attacker with perfect eavesdropping capabilities, where an at-tacker receives all beacon messages sent over the network. They assume that vehicles broadcast their location and velocity at regular intervals, but with pseudonyms that change for every packet to completely anonymize the transmissions. Their experiments showed that at high beaconing rates of a beacon a second or faster and less than 100 vehicles, they were able to track vehicles for on average 800 out of the 1000 seconds in their simulations. Increasing the vehicle density to be between 100 and 250 vehicles saw the average tracking time drop to 700 seconds. A beaconing rate of a beacon every two seconds sees an average of less than 400 seconds of tracking when there are even 50 vehicles, and a drop to 150 seconds when there are 100 vehicles. Beaconing rates slower than a beacon every two seconds only saw tracking of 100 seconds when there were less than 25 vehicles in the system, and there was no substantial tracking after 50 vehicles. Their attack is also dependent on accurate position information. When they introduce random noise into the gathered position information, a random o set anywhere between one to five meters decreased tracking by 200 down to 700 seconds. They also explore the effectiveness of their tracking with varying equipment rates. They actually see an increase in average tracking time, up to nearly 900 seconds when the equipment rate is 10 or 20%. This is due the much smaller number of vehicles being tracked, as they can only track equipped vehicles, and lower equipment rates mean the likelihood of equipped vehicles crossing paths is less likely.

Conclusion

In conclusion, expanded applications in VANETs have brought into question vehicular privacy concerns. The use of pseudonyms is the key means of defending user privacy when broadcasting messages. Pseudonym change strategies are algorithms used to determine when vehicles should change pseudonyms, and research into different change strategies have aimed to maximize the privacy gains while minimizing the number of pseudonym changes. Research has also been done into different potential attacks that an attacker can use to undermine vehicular privacy. Researching both defenses and at-tacks are important to defending vehicular privacy. The development of new attacks can be used to test the effectiveness of defenses against such attacks.

References

[1] M. L. Sichitiu and M. Kihl, \Inter-vehicle communication systems: a survey," IEEE Communications Surveys & Tutorials, vol. 10, no. 2, 2008.

[2] I. Bilogrevic, I. Aad, P. Ginzboorg, V. Niemi, J.-P. Hubaux et al., \Track me if you can: On the e ectiveness of context-based identi er changes in deployed mobile networks," in Proceedings of the 19th Annual Network & Distributed System Security Symposium (NDSS 2012), no. EPFL-CONF-169827. Internet Society, 2012.

[3] A. Musa and J. Eriksson, \Tracking unmodi ed smartphones using wi-monitors," in Proceedings of the 10th ACM conference on embedded network sensor systems. ACM, 2012, pp. 281{294.

[4] J. Liao and J. Li, \E ectively changing pseudonyms for privacy protec-tion in vanets," in Pervasive Systems, Algorithms, and Networks (IS-PAN), 2009 10th International Symposium on. IEEE, 2009, pp. 648{ 652.

[5] M. Raya, P. Papadimitratos, and J.-P. Hubaux, \Securing vehicular communications," IEEE wireless communications, vol. 13, no. 5, 2006.

[6] P. Papadimitratos, L. Buttyan, J.-P. Hubaux, F. Kargl, A. Kung, and M. Raya, \Architecture for secure and private vehicular communica-tions," in Telecommunications, 2007. ITST'07. 7th International Con-ference on ITS. IEEE, 2007, pp. 1{6.

[7] R. M. Ishtiaq Roufa, H. Mustafaa, S. O. Travis Taylora, W. Xua, M. Gruteserb, W. Trappeb, and I. Seskarb, \Security and privacy vul-nerabilities of in-car wireless networks: A tire pressure monitoring sys-tem case study," in 19th USENIX Security Symposium, Washington DC, 2010, pp. 11{13.

[8] J. Petit, F. Schaub, M. Feiri, and F. Kargl, \Pseudonym schemes in ve-hicular networks: A survey," IEEE communications surveys & tutorials, vol. 17, no. 1, pp. 228{255, 2015.

[9] N. Plewtong, \Modeling adversarial insider vehicle in mix zones," Mas-ter's thesis, California Polytechnic State University San Luis Obispo, San Luis Obispo CA 93407, 3 2018.

[10] J.-H. Song, V. W. Wong, and V. C. Leung, \Wireless location privacy protection in vehicular ad-hoc networks," Mobile Networks and Appli-cations, vol. 15, no. 1, pp. 160{171, 2010.

[11] M. Gerlach and F. Guttler, \Privacy in vanets using changing pseudonyms-ideal and real," in Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th. IEEE, 2007, pp. 2521{2525.

[12] D. Da Silva, T. Ann Kosa, S. Marsh, and K. El-Khatib, \Examining privacy in vehicular ad-hoc networks," in Proceedings of the second ACM international symposium on Design and analysis of intelligent vehicular networks and applications. ACM, 2012, pp. 105{110.

[13] K. Sampigethaya, M. Li, L. Huang, and R. Poovendran, \Amoeba: Ro-bust location privacy scheme for vanet," IEEE Journal on Selected Areas in communications, vol. 25, no. 8, 2007.

[14] D. F•orster, F. Kargl, and H. L•ohr, \A framework for evaluating pseudonym strategies in vehicular ad-hoc networks," in Proceedings of the 8th ACM Conference on Security & Privacy in Wireless and Mobile Networks. ACM, 2015, p. 19.

[15] A. R. Beresford and F. Stajano, \Mix zones: User privacy in location-aware services," in Pervasive Computing and Communications Work-shops, 2004. Proceedings of the Second IEEE Annual Conference on. IEEE, 2004, pp. 127{131.

[16] B. Wiedersheim, Z. Ma, F. Kargl, and P. Papadimitratos, \Privacy in inter-vehicular networks: Why simple pseudonym change is not enough," in Wireless On-demand Network Systems and Services (WONS), 2010 Seventh International Conference on. IEEE, 2010, pp. 176{183.

[17] R. E. Kalman, \A new approach to linear ltering and prediction prob-lems," Journal of basic Engineering, vol. 82, no. 1, pp. 35{45, 1960.