Title

Privacy Risks in Vehicle Grids and Autonomous Cars

Permalink

https://escholarship.org/uc/item/0j2088kb

Authors

Joy, Joshua
Gerla, Mario

Publication Date

2017-10-01

Peer reviewed

Privacy Risks in Vehicle Grids and Autonomous Cars

Joshua Joy
University of California - Los Angeles
jjoy@cs.ucla.eduom

Mario Gerla
University of California - Los Angeles
gerla@cs.ucla.edu

ABSTRACT

Traditionally, the vehicle has been the extension of the manual ambulatory system, docile to the drivers' commands. Recent advances in communications, controls and embedded systems have changed this model, paving the way to the Intelligent Vehicle Grid. The car is now a formidable sensor platform, absorbing information from the environment, from other cars (and from the driver) and feeding it to other cars and infrastructure to assist in safe navigation, pollution control and traffic management. The next step in this evolution is just around the corner: the Internet of Autonomous Vehicles. Like other important instantiations of the Internet of Things (e.g., the smart building, etc.), the Internet of Vehicles will not only upload data to the Internet with V2I. It will also use V2V communications, storage, intelligence, and learning capabilities to anticipate the customers' intentions and learn from other peers. V2I and V2V are essential to the autonomous vehicle, but carry the risk of attacks. This paper will address the privacy attacks to which vehicles are exposed when they upload private data to Internet Servers. It will also outline efficient methods to preserve privacy.

1 INTRODUCTION

The urban fleet of vehicles is evolving from a collection of sensor platforms that provide information to drivers and upload filtered sensor data (e.g., GPS location, road conditions, etc.) to Internet Servers; to a network of autonomous vehicles that exchange their sensor inputs among each other in order to optimize several different utility functions. One such function, and probably the most important for autonomous vehicles, is prompt delivery of the passengers to destination with maximum safety and comfort and minimum impact on the environment. We are witnessing today in the vehicle fleet the same evolution that occurred ten years ago in the sensor domain from Sensor Web (i.e., sensors are accessible from the Internet to get their data) to Internet of Things (the computers with embedded sensors are networked with each other and make intelligent use of the sensors). In the intelligent home, the IOT formed by the myriad of sensors and actuators that cover the house internally and externally, can manage all the utilities in the most economical way, with maximum comfort to residents and virtually no human intervention. Similarly, in the modern energy grid, the IOT consisting of all components large and small can manage

power loads in a safe and efficient manner, with the operators now playing the role of observers. In the vehicular grid, the Internet of Vehicles (IOV) is more complex than the smart home and smart energy grid IOTs. In fact there are many different "Things" in the IOV. Namely:

(1) External sensors (GPS, cameras, lidars etc)
(2) Internal automotive sensors and actuators (brakes, steering wheel, accelerator, etc)
(3) Internal cockpit sensors (driver's state of health, alertness, tone of voice, health sensors like the Ford heart monitor seat, etc)
(4) The Driver's messages (tweets, Facebook, other crowd-sourced info, etc) are also measurable sensor outputs that characterize the state of the system and of the driver.
(5) Vehicle's beacons, alarms report on the Vehicle state; say, position, key internal parameters, possible dangers, etc.

This complex picture (of sensors and stakeholders) tells us that IOVs are different from other IOTs. What sets them apart from other IOTs are the following properties/characteristics:

(1) **Mobility**:
    (a) IoVs Must manage mobility and wireless bottleneck
    (b) They must guarantee motion privacy
(2) **Safety critical Applications**
    (a) This implies low latency requirements
(3) **V2V**:
    (a) V2V is critical for safety, low latency apps (eg, platoons)
(4) **Attacks**:
    (a) Security and DDoS attacks (from hackers and form malicious agents) are made possible by V2V.

In the vehicular network, like in all the other IOTs, when the human control is removed, the autonomous vehicles must efficiently cooperate to maintain smooth traffic flow in roads and highways. Visionaries predict that the self-driving vehicles will behave much better than human drivers, handling more traffic with lower delays, less pollution and better driver and passenger comfort. However, the complexity of the distributed control of hundreds of thousands of cars cannot be taken lightly. If a natural catastrophe suddenly happens, say an earthquake, the vehicles must be able to coordinate the evacuation of critical areas in a rapid and orderly manner. This requires the ability to efficiently communicate with each other and also to discover where the needed resources are (e.g., ambulances, police vehicles, information about escape routes, images about damage that must be avoided, etc.). Moreover, the communications must be secure, to prevent malicious attacks that in the case of autonomous vehicles could be literally deadly since there is no standby control and split second chance of intervention by the driver (who meantime may be surfing the web).

All of these functions, from efficient communications to distributed processing over various entities, will be provided by an

emerging compute, communications and storage platform specifically designed for vehicles—the *Vehicular Cloud*. The Vehicular Cloud is justified by several observed trends:

(1) Vehicles are becoming powerful sensor platforms
    (a) GPS, video cameras, pollution, radars, acoustic, etc
(2) Spectrum is becoming scarce => Internet upload of all the sensor outputs expensive and besides infeasible
(3) More data is cooperatively processed by vehicles rather than uploaded to Internet:
    (a) road alarms (pedestrian crossing, electr. brake lights), platoon coordination signals, intersection announcement, etc
(4) Distributed Surveillance (video, mechanical, chemical sensors)
    (a) Must be locally supported, deployed
(5) Protection from DDoS attacks, must be done locally, via the Vehicular Cloud

To support the above functions, the mobile Vehicle Cloud provides several basic services, from routing to content search, through standard, open interfaces that are shared by all auto manufacturers.

2 EMERGING APPLICATIONS

A number of applications have emerged in recent years, leveraging V2I and V2V

- Safe Navigation
- Crash prevention; platoon stability; shockwaves
- Content Download/Upload
- News, entertainment, location relevant info download; ICN
- Video upload (eg remote drive, Pic-on-wheels, accident scene, etc)
- Sensor Data gathering
- Forensics; driver behavior; traffic crowdsource; ICN
- Privacy preserving data analysis
- Intelligent Transport
- Efficient routing to mitigate congestion/pollution
- Vehicle Autonomy
- Autonomous, self driving vehicles, etc

We will focus on the intelligent transport application, because it is an application that leverages both V2V and V2I communications. For this application, we will highlight the security and privacy risks caused by vehicular communications.

2.1 Intelligent navigation - from Dash Express to WAZE

Dash Express revolutionized the navigator business in 2008 by exploiting Time and Speed crowdsensed by its customers []. Namely, cars periodically submit Time and Speed reports. Using customers reports, an uptodate map of road delays is computed and accurate navigation instructions are dispatched to cars. Current Navigators are mostly based on the same crowdsourcing model. For example WAZE (by Google) is driven by customer reports supplied to the Server in the Cloud. It is accessed via V2I (DSRC, WIFI or LTE).

The Centralized, Cloud based Navigator Server has access to many other services in the Cloud (eg, historic traffic data, road repair schedules, driver habits, metereology data, pollution data, e-vehicle recharge station locations, etc). These data sources enable

the support of many advanced features like: optimization of routes; Minimization of pollution (eco routing); Traffic flow balancing; Arrival time control on preferred routes; Combined traffic and congestion control; increase patrolling to discourage bad driving behavior. However, centralized traffic management alone cannot react promptly to local traffic perturbations (WAZE has a reaction time of 10-15min). For example, a doubled parked truck in the next block; a recent traffic accident; a sudden queue of traffic on the preplanned route forces the driver to wait up to 15 min before WAZE discover the traffic blocking and find an alternate path. Namely, for scalability reasons, the Internet based Navigator Server cannot micro-manage traffic

This is where distributed traffic management can come to help! In fact, it was shown that the distributed approach is a good complement to centralized supervision, Leontiadis et. al. [10]. In the referenced paper the distributed, totally crowdsourced scheme "CATE: Comp Assisted Travel Environment" is introduced. In a nutshell, the vehicles crowd source traffic information and build traffic load data base:

(1) estimate traffic from own travel time;
(2) share it with neighboring vehicles (with V2V in an ad hoc manner)
(3) dynamically recompute the best route to destination

Interestingly, both Centralized and Distributed navigation systems lead to security issues. In the past we have investigated the vulnerability of the distributed V2V scheme to BOTNET attacks launched by compromised cars [xyz]. The compromised cars manage to propagate false information, via V2V and lure honest cars in a major traffic bottleneck in a couple of minutes [8]! The Centralized Navigator protects from BOTNETs, but exposes customers to Privacy attacks, as described below.

2.2 Security Problem: Privacy violations in V2I communications

The Centralized Navigators also have security problems and can lead to Communication Privacy Violations. In fact, with centralized navigators, Cars upload their position, velocity and intended destination to the Navigator. For example:

- WAZE delivers vehicle position and traffic conditions to GOOGLE traffic
- UBER vehicles upload passenger and vehicle status to UBER Server
- LTE providers can trilaterate and localize the vehicles as they connect to the Internet

The collected data can be used by the Navigation servers to track users and discover their driving habits, favorite hot spots, etc. Naturally, Service Providers like GOOGLE, UBER and Cellular Companies are committed to protect customer privacy. However, privacy guarantees have been often broken in the past (intentionally or by mistake). In the Waze Privacy Attack, Waze allows remote customers to view current traffic in an arbitrary window. By moving the window, the attacker tracks the victim. In the Waze DDoS Attack, the malicious customer impersonates multiple WAZE vehicles in a small area, simulating traffic bottleneck [12].

In the remainder of this paper, we focus on the Privacy violation issue, a problem common to all applications that upload mobile data from IOT or IOV to Servers in the Cloud. We formally define the problem, introduce an efficient, scalable solution, Haystack, and evaluate it on a real distribution of driver habit responses collected by Triple AAA and stored in the publicly accessible Safety Culture Index [1]

3 RELATED WORK

Differential privacy [3–6] has been proposed as a mechanism to privately share data such that anything that can be learned if a particular data owner is included in the database can also be learned if the particular data owner is not included in the database. To achieve this privacy guarantee, differential privacy mandates that only a sublinear number of queries have access to the database and that noise proportional to the global sensitivity of the counting query is added (independent of the number of data owners).

The randomized response based policies [7, 9, 11, 13] satisfies the differential privacy mechanism as well as stronger mechanisms such as zero-knowledge privacy. However, the accuracy of the randomized response mechanism quickly degrades unless the coin toss values are configured to large values (e.g., greater than 80%).

4 HAYSTACK PRIVACY

We now introduce the Haystack Privacy mechanism. Our goal is for scalable privacy whereby as more individuals participate the privacy guarantee becomes stronger while simultaneously we would like to maintain constant error in the worst case. We motivate the need for scalable privacy with the following example.

Suppose we issue a counting query whereby we are interested in how many human drivers aggressively accelerate and tailgate another vehicle. Say the query is only targeted at those aggressive drivers and 95 out of 100 queried drivers tailgate over vehicles. Clearly if an adversary knows that a particular driver participated in the study there no privacy as any adversary is able to guess with a greater than 90% success rate.

Is there a way to address this privacy breach? One possible solution is to query a larger population, say an entire city regardless of prior knowledge of their driver behavior. This would allow us to collect results from a more diverse population. Say we queried 1 million drivers while only 95 of the drivers are aggressive and tailgate. Now to perform a privacy breach an adversary must determine which 95 drivers of the total population are the aggressive drivers.

We can go further. The population can be diversified while we simultaneously learn more information. For example, we could query for tailgaters, excessive lane changing, speeding, and normal behavior. Querying additional attributes accomplishes two key properties. First, we are able to increase the participating population. Second, we are able to learn additional features such as counts of those that speed and do excessive lane changing.

However, there is a concern. Data owners that do not truthfully respond "Yes" they are a tailgater add privacy protection at the cost of distorting the underlying distribution. In our example, the distribution changes from 95% of the population to less than 0.01% of the population. The question then becomes can we preserve any notion of accuracy?

Performing sampling over the "distorted" distribution with a large population of say one million will incur a large sampling error that will dominate the estimation. Instead, we run a multi-round protocol and fix certain coin tosses across rounds in order to eliminate the error due to sampling. This way, we maintain constant error.

We now describe our Haystack Privacy Mechanism in detail.

4.1 Haystack Privacy Mechanism

Illustration. To illustrate and demonstrate the mechanism, we employ the following example. Suppose we are interested in the distribution of the degree of speeding behavior. Aggressive driving is a factor for more than one-half of all traffic fatalities and speeding is a factor in one-third of all fatal crashes [2]. Recent studies have shown that more than half of all drivers surveyed admitted to speeding more than 15% of the posted speed limit in the past 30 days [2].

Suppose a data owner was speeding. First, the data owner should discretize the amount they were speeding by. Suppose we discretize the speeds below and exceeding the posted speed limit as follows. Below the speed limit are as follows " $-15 \sim -11$" is group 1, " $-10 \sim -6$" is group 2, " $-5 \sim -0$" is group 3. Exceeding the speed limit is "$1 \sim 5$" is group 4, "$6 \sim 10$" is group 5, "$11 \sim 15$" is group 6 etc.

Say the data owner was speeding by 10 mph. Then the data owner discretizes their speed integer 5 value (group 5).

In the first round, the data owner tosses a multi-sided die. One side samples whether the data owner should respond truthfully for their location ID. The remaining sides selects a group ID for the data owner to respond. Say the number of speeding groups is G. We will use a total of $G + 1$ groups as we will see below.

Suppose in the first round the data owner is sampled and selected. The data owner should respond "Yes" to group 0 (used as a calibration step). The remaining data owners also toss a multi-sided die and respond with the group number corresponding to the die number. A privatized sum is computed by aggregating the "Yes" counts in each group.

In the second round the sampled data owner should respond truthfully for their given group ID. The remaining data owners stay with their first round responses. A privatized sum is computed by aggregating the "Yes" counts in each group.

Finally, population size estimation is carried out for each group ID by subtracting the privatized sum in round two from round one and dividing by the sampling parameter.

The following three privacy observations are made. First, a majority of the population provides privacy noise by randomly responding either "Yes" or "No" regardless of their truthful response. Second, plausible deniability is provided as each data owner probabilistically responds opposite of their truthful response. Finally, *every* data owner acts as a potential candidate for the truthful population. Our assumption is that every data owner is active in both rounds and only the aggregate counts are released.
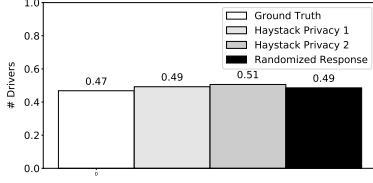
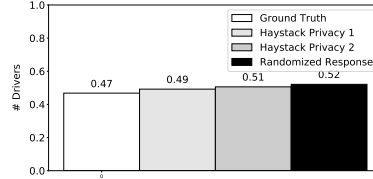**Figure 1: (Excessive Speeding) 3,896 respondents.**

Figure 2: (Excessive Speeding) 100,000 respondents (3,896 respondents and 96,104 add chaff).
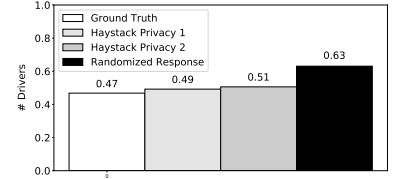

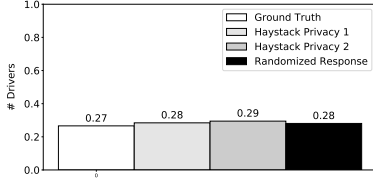**Figure 3: (Excessive Speeding) 1,000,000 respondents (3,896 respondents and 996,104 add chaff).**

Figure 4: (Texting and Driving) 3,896 respondents.
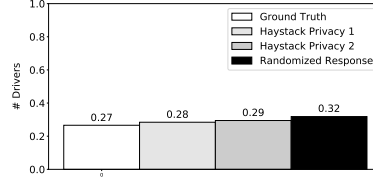

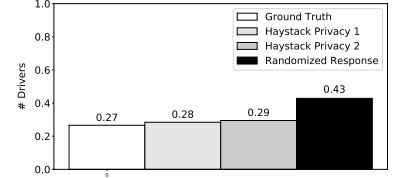**Figure 5: (Texting and Driving) 100,000 respondents (3,896 respondents and 96,104 add chaff).**

Figure 6: (Texting and Driving) 1,000,000 respondents (3,896 respondents and 996,104 add chaff).

Haystack Privacy Mechanism. We now formally describe the Haystack Privacy mechanism.

(Round One) Let G represent all outputs for which the data owner does not truthfully respond "Yes". Let G' be the special output for which the data owner truthfully responds "Yes".

In the first round the sampled and selected data owner responds to $\varnothing$ (corresponding to group 0). The remaining data owners respond according to toss of the multi-sided die.

$$Round\ One_{Yes} = \begin{cases} \varnothing & \textbf{with probability } \pi_s \\ G, G' & \text{with probability } \pi_G \end{cases} \quad (1)$$

That is, all sampled and selected data owners respond to $\varnothing$ while the remaining data owners randomly respond.

(Round Two) In the second round the sampled and selected data owner now responds with their truthful response. The remaining data owners stay with their round one response.

$$Round\ Two = \begin{cases} \varnothing & \textbf{with probability } 0 \\ \textbf{G}' & \textbf{with probability } \pi_s \\ G, G' & \text{with probability } \pi_G \end{cases} \quad (2)$$

That is, all the sampled and selected data owners respond truthfully allowing us to now perform the estimation.

(Expected Values) Let G_1 be a given group value in the first round. For a given value of G and G' let Yes_{pop} refer to the truthful "Yes" fraction of the population and No_{pop} refer to the truthful "No"

fraction of the population. The entire population is represented by $TOTAL$. The first round of expected values are as follows.

$$E[G_1] = \pi_G \times TOTAL \quad (3)$$

That is, for each value both populations randomly contribute.

The second round the expected values now include the sampled population.

$$E[G_2] = \pi_G \times TOTAL + \pi_s \times Yes_{pop} \quad (4)$$

That is, everyone randomly contributes. The sampled and selected percentage truthfully respond.

(Estimator) To solve for the YES population we subtract the first round from the second round and repeat for each output value as follows:

$$YES = \frac{Private\ Sum_{G,2} - Private\ Sum_{G,1}}{\pi_s} \quad (5)$$

The sampled and selected population, by not participating in round one, allows us to baseline the privacy noise and perform estimation for the sampled truthful population.

5 EVALUATION

We evaluate the Haystack Privacy mechanism over a distributions of inputs from real drivers obtained from Triple AAA Safety Culture Index [1]. We assign virtual identities to each respondent. The population is a random sample of 3,896 U.S. residents of driving age.

Figures 1, 2, 3 shows the results from the first experiment, namely, the distribution of drivers that have excessive speeding greater than 10mph over the posted speed limit in the past 30 days. We increase the number of drivers not participating in the particular study and who do not exhibit excessive speed. We show the scaling effects. Upper bounds are shown with a 95% confidence interval. The coin toss probabilities are fixed as follows. Haystack Privacy 1 $\pi_s = 0.45$ and Haystack Privacy 2 $\pi_s = 0.25$. Randomized Response $flip_1 = 0.8$ and $flip_2 = 0.2$.

Figures 4, 5, 6 shows the results from the second experiment, namely the distribution of drivers that have written and sent texts while driving within the past 30 days. Like in experiment 1, we increase the number of drivers not participating in the particular study and who do not write/send text and show the scaling effects. Upper bounds are shown with a 95% confidence interval. The coin toss probabilities are fixed as follows. Haystack Privacy 1 $\pi_s = 0.45$ and Haystack Privacy 2 $\pi_s = 0.25$. Randomized Response $flip_1 = 0.8$ and $flip_2 = 0.2$.

We compare Haystack Privacy to the conventional Randomized Response based on two biased coins' tossing. The Haystack Privacy mechanism maintains constant error while the Randomized Response accrues error as the population scales. It should also be noted that as the underlying distribution to estimate tends to smaller values, the Randomized Response has difficulties in performing accurate estimation. This has implications for large and diverse datasets whereby we are interested in estimating and understanding non-frequently occurring phenomena (e.g., hard to explain traffic accidents and environmental factors).

6 CONCLUSION

In this paper we have demonstrated that data can be privately collected into a common open data vehicular database to be shared amongst multiple collaborators. We introduce the concept of Haystack Privacy, which scales well by increasing the privacy strength as more data owners participate yet maintaining accuracy. Haystack Privacy easily outperforms Randomized Response. We believe this is a new direction in open data vehicular research.

REFERENCES

[1] 2012 Traffic Safety Culture Index [n. d.]. 2012 Traffic Safety Culture Index. https://www.aaafoundation.org/sites/default/files/2012TrafficSafetyCultureIndex.pdf. ([n. d.]). https://www.aaafoundation.org/sites/default/files/2012TrafficSafetyCultureIndex.pdf

[2] Aggressive Driving [n. d.]. Aggressive Driving . https://www.aaafoundation.org/aggressive-driving. ([n. d.]). https://www.aaafoundation.org/aggressive-driving

[3] Cynthia Dwork. 2006. Differential Privacy. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II (Lecture Notes in Computer Science)*, Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener (Eds.), Vol. 4052. Springer, 1–12. https://doi.org/10.1007/11787006_1

[4] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. 2006. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *EUROCRYPT*.

[5] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC*.

[6] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3-4 (2014), 211–407. https://doi.org/10.1561/0400000042

[7] James Alan Fox and Paul E Tracy. 1986. *Randomized response: a method for sensitive surveys*. Beverly Hills California Sage Publications.

[8] Mevlut Turker Garip, Peter Reiher, and Mario Gerla. 2016. Ghost: Concealing vehicular botnet communication in the VANET control channel. In *2016 International Wireless Communications and Mobile Computing Conference (IWCMC), Paphos, Cyprus, September 5-9, 2016*, Mario Gerla, George C. Hadjichristofi, and Mohsen Guizani (Eds.). IEEE, 1–6. https://doi.org/10.1109/IWCMC.2016.7577024

[9] Bernard G Greenberg, Abdel-Latif A Abul-Ela, Walt R Simmons, and Daniel G Horvitz. 1969. The unrelated question randomized response model: Theoretical framework. *J. Amer. Statist. Assoc.* 64, 326 (1969), 520–539.

[10] Ilias Leontiadis, Gustavo Marfia, David Mack, Giovanni Pau, Cecilia Mascolo, and Mario Gerla. 2011. On the Effectiveness of an Opportunistic Traffic Management System for Vehicular Networks. *IEEE Trans. Intelligent Transportation Systems* 12, 4 (2011), 1537–1548. https://doi.org/10.1109/TITS.2011.2161469

[11] Ajit C. Tamhane. 1981. Randomized Response Techniques for Multiple Sensitive Attributes. *J. Amer. Statist. Assoc.* 76, 376 (1981), 916–923. https://doi.org/10.1080/01621459.1981.10477741

[12] Gang Wang, Bolun Wang, Tianyi Wang, Ana Nika, Haitao Zheng, and Ben Y. Zhao. 2016. Defending against Sybil Devices in Crowdsourced Mapping Services. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services, MobiSys 2016, Singapore, June 26-30, 2016*, Rajesh Krishna Balan, Archan Misra, Sharad Agarwal, and Cecilia Mascolo (Eds.). ACM, 179–191. https://doi.org/10.1145/2906388.2906420

[13] Stanley L Warner. 1965. Randomized response: A survey technique for eliminating evasive answer bias. *J. Amer. Statist. Assoc.* 60, 309 (1965), 63–69.