Predictive Cruise Control with Private Vehicle-to-Vehicle Communication for Improving Fuel Consumption and Emissions

Xueru Zhang[1*], Chunan Huang[2*], Mingyan Liu[1], Anna Stefanopoulou[2], Tulga Ersal[2]
[1]Department of Electrical and Computer Engineering, University of Michigan
[2]Department of Mechanical Engineering, University of Michigan
[*]Equal contribution

Future traffic information through vehicular communication allows connected and automated vehicles to optimize their speed trajectories and drive more safely and efficiently through predictive controllers. Sharing accurate information about the vehicle allows such controllers perform best, but may raise privacy concerns. To improve privacy guarantee over the shared information while preserving its utility for predictive controllers, this paper proposes a novel information perturbation mechanism, as opposed to the baseline of independently perturbing the data in each broadcast. Specifically, the mechanism is applied to the transmitted vehicle speed, and this perturbed data is used in an optimal speed planner to design a fuel and emissions efficient speed trajectory. Results show a deterioration of the controller performance when privacy is taken into consideration under the baseline method. With the proposed method, the controller performance is improved while providing the same privacy guarantee. It is shown that controller design is also affected by the choice of perturbation mechanism.

Index Terms—vehicular communication, differential privacy, connected automated vehicle, predictive control

I. INTRODUCTION

A vehicle's movement on the road is constrained by the route, road and traffic conditions it encounters, such as the motion of the neighboring vehicles, traffic signals, and local road and weather conditions. Knowing this information and future movements of the surrounding vehicles would allow an automated vehicle to drive more efficiently. In car-following scenarios, knowing the future speed profile of the leader vehicle has been shown to be beneficial for an automated follower vehicle to drive more safely and efficiently [1, 2], as predictive speed controllers can use this information to design an optimal trajectory [2, 3].

Regardless of the optimization objective, most of the literature on optimal speed planning using predictive controllers for connected and automated vehicles assumes that the future information is available with high accuracy, either obtained directly from the leader vehicle, or inferred from vehicle-to-vehicle (V2V) communication with accurate information about the leader vehicle [2, 4]. In most real traffic scenarios, a vehicle's speed is hard to predict accurately due to uncontrollable factors. Researchers have shown that inaccurate speed prediction in predictive controllers may increase both the risk of collision and fuel consumption compared to the case with accurate information [4]. To address this problem, they propose two stochastic model predictive controllers and show that with a certain speed prediction, better performances are achieved with these stochastic controllers than the deterministic one [4]. However, with prediction error, none of these controllers recover the performance under accurate information.

Even if not knowing the true information may degrade performance, it is not practical to assume that the true information is available. One potential reason is the drivers' privacy concerns when their vehicular information (e.g., speed, location) is transmitted to other, untrusted vehicles. Examples of these privacy concerns include: (1) Tracking and stalking: Locations along with other publicly available information can be used to identify a driver's personal information, thereby enabling stalking. Research shows that 5% of U.S. workers can be uniquely identified by just knowing their home and work areas [5]. (2) Traffic enforcement: Drivers may be concerned that V2V tracking could facilitate automated issuance of traffic citations. If certain privacy guarantees can be provided, however, drivers may be more willing to share their personal driving information. Government organizations also acknowledge the need to address privacy before implementing V2V communication technologies [6].

Various notions of privacy have been suggested for applications in vehicular networks. They can be roughly classified into anonymity-based and perturbation-based methods. The former de-identifies each vehicle to provide privacy, either by replacing the real unique identifier of each vehicle with some variable and temporary pseudonyms, or by adopting the k-anonymity technique [7], where at least k vehicles would share the same set of attributes (that are indirectly related to identifiers) and form an anonymity set; vehicles within the same set cannot be distinguished from each other. Solely changing the pseudonym cannot protect vehicles from being tracked over time [8], and combining pseudonym with k-anonymity is therefore preferred [9].

Perturbation-based methods provide privacy by perturbing transmitted/shared information. Typically, the notion of differential privacy [10] is adopted, under which each vehicle transmits a noisy version of its actual information. Different from k-anonymity, where an adversary's background knowledge
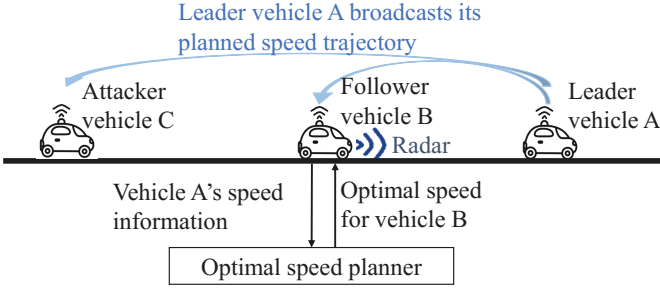
Fig. 1. Traffic scenario considered in this paper.

must be predefined, differential privacy can protect against adversaries with any side information and is much stronger. Researchers have applied differential privacy to vehicular networks [11, 12]. Most algorithms are either only applicable to a specific application (e.g., binary classification [12]) or to a setting with a trustworthy, third-party data collector gathering speed/location data from all vehicles [11]. Therefore, designing a more general differentially private method that can release any type of data in real-time is of interest.

This study considers a car-following scenario, where a leader vehicle generates its speed profile with the differential privacy constraint and broadcasts it periodically. Based on this information, the follower vehicle designs its own optimal speed using a planner that aims at reducing fuel consumption and tailpipe emissions. To satisfy the differential privacy constraint for the leader while achieving a sufficiently accurate speed preview for the follower, a new perturbation method is introduced to balance the trade-off between privacy and accuracy. Simulation results show that the proposed method can generate differentially-private speed broadcasts with sufficiently accurate information for improving fuel and emissions performance through predictive speed planning, while with the baseline method, the perturbed information with the same differential privacy guarantee is nearly useless for improving these performances.

We next present the problem formulation and definition of differential privacy followed by the two different perturbation mechanisms.

II. COMMUNICATION-AIDED SPEED PLANNING

A. Problem formulation

We focus on the application of predictive speed planning through private V2V communication. The traffic scenario considered is described in Fig. 1. We assume that vehicle A is broadcasting its information, and other vehicles within the communication range of vehicle A, such as vehicles B and C, can receive the broadcast information. The information sent is vehicle A's predicted speed in the next few seconds to a minute.

The vehicle that is immediately following vehicle A, e.g. the follower vehicle B in Fig. 1, can use the leader's information for predictive speed planning. The follower is equipped with a communication receiving unit, an on-board optimal speed controller, and a radar to measure the current speed and the

inter-vehicular distance to the leader vehicle A. The optimal speed controller integrates vehicle A's future speed received through V2V communication and radar measurements to estimate the motion of A in the near future, and uses it as a traffic constraint for vehicle B to optimize its future speed trajectory. For instance, if B knows that A is going to perform an acceleration followed by a deceleration and an extended stop, then B can utilize this information to determine the best way to drive in terms of fuel economy, driving comfort, and/or emissions through the optimal speed planner in Fig. 1.

However, as mentioned in Sec. I, drivers' personal information can further be inferred from the vehicles' speeds/locations. Thus, privacy concerns inevitably arise when the vehicles' information is disseminated among connected vehicles. In the car-following scenario, the leader's true speed is eventually revealed to the follower completely due to the radar. The driver of the leader may nonetheless wish to keep their information private from other non-follower vehicles or roadside units that are within the communication range of the leader and can receive the leader's future speed trajectory, but cannot detect the actual speed directly with radar. An example is vehicle C in Fig. 1 that is referred to as attacker vehicle.

If the precise future speed information of vehicle A is transmitted to vehicle C, then the information such as whether vehicle A is speeding or not, whether its driver is erratic or not, etc., will be revealed directly. To keep vehicle A's speed private from the attacker vehicle C, A should broadcast private versions of the future speed profile instead.

We adopt the perturbation-based method by adding noise to trajectories and use differential privacy as a notion of privacy to measure the privacy risk of each vehicle. We assume the attacker can have any side information about the leader and follower vehicles, including all algorithms they implement, the noise distributions the leader uses, etc. The attacker vehicles, by receiving the same noisy trajectory broadcasts as the follower, can extract almost the same amount of information from these trajectories about the leader as the follower. We say almost, because the follower has additional information from the radar while the attackers do not. Therefore, to preserve A's privacy from C, and simultaneously provide useful information to B's optimal speed planner, the perturbation should be carefully designed to balance the trade-off between privacy and accuracy.

B. Differential privacy [10]

Differential privacy centers around the idea that the output of a certain mechanism or computational procedure should be statistically similar given singular changes to the input, thereby preventing meaningful inference from observing the output.

In our V2V communication context, this differential privacy is provided by perturbing the true information coming from each vehicle. Specifically, each vehicle transmits a noisy version of its future speed trajectory, by either randomizing the trajectory directly or randomizing the algorithm that generates the trajectory. After randomization, the output (trajectory) is no longer deterministic but a random vector with a certain distribution that depends on the form of the added noise (e.g., Gaussian, Laplace, etc.).
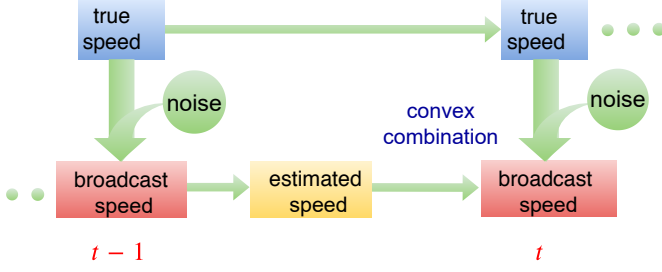
Fig. 2. Two-step illustration of the proposed method: Vehicle A's perturbed speed at each broadcast is determined by the convex combination of the true speed and the estimation from the previous broadcast.


Fig. 3. (Upper) Transmitted information perturbed by the baseline method and the proposed method, and (lower) information available to vehicle B at two consecutive time steps.

When done correctly, the conditional probability distributions of this random output given all possible speed trajectories should be similar, and differential privacy guarantees this holds for any possible random output. The bound on the differences of those distributions can be used to quantify the privacy risk: a smaller discrepancy implies stronger privacy. Specifically, this difference is characterized by the log-likelihood ratio; thus the privacy risk can vary from zero to infinity. With a small privacy risk, an attacker has no confidence in guessing the true speed when given a noisy speed, as all values after perturbation are likely to give the same output, with similar likelihood values. Consider an extreme case where the distributions are the same; then the inference that an attacker can make from the transmitted output will be the same, regardless of its true value, thereby conferring complete privacy protection (zero privacy risk). In contrast, if the precise trajectory is transmitted without any perturbation, then the probability of observing this output would be 1 for this particular trajectory and 0 for all other possible trajectories, resulting in infinite privacy risk.

Differential privacy is a worse-case measure; i.e., the bound is over all possible random outputs and all possible inputs. It is a strong guarantee, as it can protect against attackers with any side information. Moreover, it is immune to post-processing; i.e., given only the differentially private output without additional information about the true data, it is impossible for attackers to make it less differentially private.

In the next section, two perturbation mechanisms are introduced to preserve differential privacy for vehicle A.

C. Two perturbation mechanisms

A naive method for vehicle A to protect its privacy is adding independent noise to the data at each broadcast, which serves as the baseline. The noise we adopt follows zero-mean Gaussian distribution. However, this method is problematic because of the temporal correlation in the data (e.g., the speed/location of vehicle A is highly correlated in consecutive broadcasts). The attacker with statistical knowledge of this correlation can be particularly hard to defend against, as it can use all transmitted information to make the inference. As a result, the privacy risk to vehicle A is accumulated over all the broadcasts and the total privacy risk can be extremely large.

To address this issue, one approach is to factor this correlation into the perturbation mechanism. We propose a new method, where for vehicle A the broadcast data in each step is based on both the broadcast data in the previous step and the true data (see Fig. 2). The idea is based on two observations: (1) Since the vehicular data generated in two consecutive broadcasts is highly correlated, and the perturbed data is also correlated with the original unperturbed data, we can use the perturbed data from the previous broadcast to estimate the true data of the current broadcast based on the statistical properties of the trajectory (e.g., mean, variance, correlation). Various estimators can be used; we adopt the commonly used Minimum Mean Square Error (MMSE) estimator. (2) Because the computation over the existing differentially private outputs will not leak additional privacy (by post-processing property), the estimation procedure does not increase the privacy risk. Thus, technically, vehicle A can broadcast just the estimates all the time. However, solely relying on estimated speed will lead to a fairly inaccurate sequence compared to the ground truth; i.e., although privacy risk does not accumulate, the estimation error does. To balance the competing needs of accuracy and privacy, we must calibrate the broadcast data using the true data. Among the potential approaches to this calibration, we simply take the convex combination of the estimate and true value as a first step. Finally, we add noise that follows zero-mean Gaussian distribution to this combination to generate the broadcast speed.

D. Privacy analysis & discussion

Note that the proposed method is a generalized version of the baseline; it reduces to the baseline if estimated speed has zero weight in the convex combination. By adjusting the weights of the estimated and true speeds, the proposed method can always improve the privacy-accuracy trade-off, and potentially significantly compared to the baseline. By repeatedly using the already transmitted speed in the estimation, less information about the real speed is revealed in each broadcast.

To guarantee the same level of privacy as the baseline, the proposed method requires less perturbation, since it reveals less information than the baseline. Thus, the proposed method has a higher accuracy under the same privacy guarantee; i.e., an improved privacy-accuracy trade-off.

Once the attacker vehicle receives private trajectory generated either by the baseline or proposed method, it may apply a noise reduction algorithm to further improve accuracy, e.g., by averaging/filtering out the random disturbance. Regardless, the privacy is unaffected by these post-processing strategies. Additionally, most noise reduction algorithms can only take a sequence of data points over multiple broadcasts as input, but not a single data in one broadcast. Thus, they do not satisfy the real-time requirement in V2V systems. In the rest of paper, we only compare the baseline and proposed method without any post-processing.

Note that the concept of sharing noisy information in V2V communication is compatible with existing standards. Cooperative Awareness Messages (CAMs) can be disseminated periodically in the Intelligent Transport Systems (ITS) network under European Telecommunications Standards Institute (ETSI) standards. Specifically, vehicular information such as speed, location and their corresponding precision with a confidence level of 95% are broadcast by each vehicle (see Annex B in ETSI EN 302 637-2). However, no precision requirement is specified. Therefore, we can adjust precision purposefully for protecting privacy. As long as the perturbation's magnitude is carefully controlled such that the information is still useful, the existing standards still apply.

E. Speed profile after perturbation

Now we present the speed profile of vehicle A after perturbation. We assume that vehicle A has its speed profile determined before starting the drive cycle, which is shown in the top plot in Fig. 3 and corresponds to the EPA Federal Test Procedure. However, this information should be perturbed when broadcast to protect privacy. Two example trajectories of the randomly perturbed drive cycle using the baseline and proposed methods are also shown in the same plot. The variances of added noise in the perturbed cycles are chosen such that their privacy guarantees are the same.

Once the trip starts, vehicle A broadcasts its perturbed future speed for every second within a time interval, the length of which is equal to the prediction horizon of the predictive speed controller on vehicle B. For illustration, the time interval is set to 10 s in the bottom plot of Fig. 3. The actual length used to obtain the simulation results in Sec. IV is 40 s. Since each speed is transmitted multiple times during multiple broadcasts, the total privacy loss will be accumulated if the noisy speed is generated independently in every broadcast. Attacker C can use time-averaging to make a better inference about the true speed. To address this issue, the same noisy speed is reused and transmitted during multiple broadcasts instead of generating a new value independently in each broadcast. The procedure is illustrated in the bottom plot of Fig. 3, where every two consecutive broadcasts, i.e., $t = 11$ s and 12 s, have an overlap, and the information during the overlapping portion

at $t = 12$ s repeats what is sent at $t = 11$ s, except for the first second of each broadcast, when true information is available to the immediate following vehicle B from the radar. Based on this, an optimal speed planner is designed and used for vehicle B as described in the next section.

III. OPTIMAL VEHICLE SPEED PLANNER

This section describes the planner we use for vehicle B that aims at optimally reducing fuel consumption and tailpipe emissions.

A. MPC formulation

A model predictive controller (MPC) is adopted as the optimal speed planner for vehicle B. The MPC decides the optimal acceleration through the following iterative process: (1) At the current time step, MPC solves an optimal control problem that minimizes a cost function over the prediction horizon subject to constraints. The cost function represents a weighted sum of fuel consumption and tailpipe NOx emissions calculated from a model. The constraints include an inter-vehicular distance constraint, which is generated from the predicted speed of the vehicle A, maximum speed and acceleration constraints, and system dynamics. (2) Even though the optimization determines the optimal acceleration trajectory for the entire prediction horizon, only the solution at the current time-step is applied to the vehicle. (3) At the next time step, steps (1) and (2) are repeated with the new information available to the optimizer.

The model used to simulate fuel consumption and tailpipe NOx emissions and the selection of cost function are described in the following subsections.

B. Fuel and emissions model

A vehicle with a diesel engine is modeled for this work. Both vehicle fuel consumption and tailpipe emissions are calculated based on knowledge of vehicle speed and acceleration, as well as air temperature, which is assumed to be constant at 25 °C. This is done by modeling vehicle longitudinal dynamics, gear shift, engine outputs (e.g., fuel consumption and inputs to aftertreatment system), aftertreatment thermal dynamics and NOx reduction ratio. The aftertreatment system comprises a Diesel Oxidization Catalyst (DOC) and a Selective Catalytic Reactor (SCR). NOx reduction process happens in the SCR and the reduction ratio is determined by the SCR temperature.

C. MPC objective function

The objective for the controller is to reduce fuel consumption and tailpipe emissions. Thus, the objective function is designed to be a weighted sum of two terms, one for fuel and the other for emissions.

Smoothing the speed trajectory leads to lower torque and power demand, which improves fuel efficiency when traveling the same distance. Thus, squared acceleration is the term in the objective function to reduce fuel consumption. On the other hand, high NOx reduction ratio is preferred to reduce tailpipe NOx consumption. NOx conversion ratio in the SCR reaches

its maximum in the range of $220-320$ °C, which, in medium-to-light duty drive cycles, corresponds to a requirement of turbine temperature staying above the threshold temperature of 240 °C. Thus, we include the squared difference between the threshold temperature and the turbine temperature if turbine temperature is lower than the threshold as the term to reduce tailpipe NOx emissions. The above two terms are used as an alternative to fuel consumption and tailpipe emissions to reduce controller complexity and computation time. It is shown in Sec. IV that this objective function is able to effectively balance fuel consumption and NOx emissions.

IV. EFFECTS OF PRIVACY ON VEHICLE PERFORMANCE

In this section, the above-described optimal speed planner is applied to vehicle B under different scenarios to assess how the perturbation employed by vehicle A affects the performance of the speed planner. For comparison, the following three scenarios are considered, where vehicle A broadcasts

- true future speed without considering privacy (Case 1).
- a private version of future speed using the proposed method (Case 2).
- a private version of future speed using the baseline method (Case 3).

We re-emphasize that speed profiles applied in Case 2 and Case 3 have the same differential privacy guarantee.

The speed planner uses a weight factor w to adjust the trade-off between minimizing acceleration (to reduce fuel consumption) and increasing turbine temperature (to reduce NOx emissions). To explore the trade-off, the weight is varied between simulations as $w = 0, 1, \ldots, 5$, where larger w means larger penalty on emissions. For all the 3 cases considered, the same set of weights is used to produce the simulation results in Fig. 4. In all the simulations vehicle A is assumed to follow the EPA Federal Test Procedure as the drive cycle.

As observed from Fig. 4, for Case 1, when true speed information of vehicle A is available to the controller, fuel consumption is reduced by 15% when the total tailpipe NOx is no more than the nominal trajectory, i.e., the case in which the follower follows the leader's true speed trajectory exactly.

Now consider the cases involving perturbations. When comparing Case 2 and Case 3 with Case 1, the overall performance worsens with decreased accuracy of the V2V information. Since the speed planner aims to minimize two competing costs, i.e., fuel and emissions, and it is impossible to achieve the minimum for both simultaneously, we evaluate the planner performance by looking at the fuel consumption when NOx emissions are the same, or by looking at the NOx emissions when the fuel consumptions are the same (Fig. 4). Graphically, the controller performance can be approximately viewed as the distance between the curve and the nominal point $(1, 1)$. Note that the trajectories in Case 2 and Case 3 have the same privacy guarantee, but Case 2 yields better performance. Hence, with the proposed method the planner performance is improved compared to the baseline without increasing the privacy risk. Note that here we assume that the desired level of differential privacy guarantee is given, and only the perturbation mechanism is designed to yield better


Fig. 4. Normalized fuel consumption vs. tailpipe NOx emissions curve for trajectories optimized over the FTP drive cycle with different weight factors w when the optimal speed planner is using speed profiles from vehicle A with different privacy guarantees. The nominal case refers to the scenario when vehicle B follows the same speed trajectory as vehicle A.
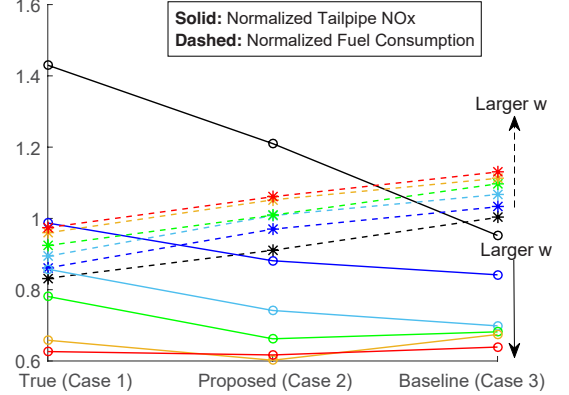

Fig. 5. Variation of normalized fuel consumption and tailpipe NOx emissions with optimal speed planners with different weight factors w when information accuracy is changed.

performance than baseline and simultaneously preserve the privacy, as we mentioned in Sec. II-C. However, as far as the authors know, it is not easy to come up with a maximum level of privacy or to design a perturbation mechanism for any privacy level, to theoretically guarantee that better performance is achieved compared with the nominal trajectory. Monte Carlo simulations can be done to approximate the maximum tolerated level for differential privacy.

The influence of prediction inaccuracy on fuel consumption and emission performances are different. As shown in Fig. 5, with the same weight factor selected in the MPC, as the information becomes more inaccurate, fuel consumption increases while tailpipe NOx emissions reacts unpredictably.

The increase in fuel consumption is expected for the following reason. When the predictive information is inaccurate, current distance and speed information from the radar on vehicle B, which is accurate and updated at every time step, will force B to accelerate or decelerate to follow vehicle A. This causes a more oscillatory drive cycle. If the information is accurate,

there is better agreement between the radar information and the V2V information, which means B is aware of vehicle A's movement in advance and can be more optimal in selecting a smoother trajectory, which costs less fuel.

For the tailpipe NOx emissions, which are affected by both the amount of engine emitted NOx and aftertreatment reaction ratio, there does not exist an as consistently monotonic relationship with the level of inaccuracy as the fuel consumption. In the simulations performed, as shown in Fig. 5, overall tailpipe NOx decreases with the level of inaccuracy when w is small, while when w is large, Case 3 creates more total tailpipe NOx than Case 2. This shows that with a small w, increase in reduction ratio is the major effect compared with the increase in engine NOx emissions under this simulated setting. However, with a large w, a higher weight is already used in the temperature related term that leads to higher aftertreatment temperature and thus enters the temperature range that produces a higher reduction ratio. Further temperature increase caused by the oscillations mentioned above does not improve the reduction ratio as much as the engine NOx increase caused by the oscillations. Thus, higher engine NOx emissions become the major effect and yield worse tailpipe NOx performance. If the control objective is maintaining the same tailpipe NOx as the nominal trajectory and reduce fuel consumption, these analyses show that the weight w in Case 2 should be larger than in Case 3. This expresses the need for an integrated design strategy, in which the tuning of the controller weight is done with consideration of the perturbation method and requirement for guaranteeing privacy.

V. Conclusion

An application of predictive speed planning in a car-following scenario is studied with differential privacy considerations. A new perturbation mechanism is proposed to guarantee a certain level of differential privacy for the leader vehicle while still providing sufficiently accurate information to the follower vehicle for speed planning with good performance. As compared to the baseline method that independently perturbs speed in every broadcast, our method generates the speed profile with the same differential privacy guarantee but with higher accuracy. The improved accuracy in the broadcast information then leads to better overall speed planning performance. Meanwhile, for more specific control objectives, selection of the control parameter is also affected by the selection of the perturbation mechanism.

The main conclusion of this work is that inaccuracies in the broadcast information of a leader vehicle that are introduced due to privacy concerns can have a significant impact on the performance of predictive speed planners the follower vehicles may utilize. This negative impact can be reduced through co-development of the differential privacy and predictive speed planning strategies. The results in this paper demonstrate the potential benefits of a more comprehensive design and analysis perspective, and motives further development of integrated strategies.

References

[1] K. C. Dey *et al.*, "A review of communication, driver characteristics, and controls aspects of cooperative adaptive cruise control (CACC)," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 2, pp. 491–509, 2016.

[2] B. Asadi *et al.*, "Predictive cruise control: Utilizing upcoming traffic signal information for improving fuel economy and reducing trip time," *IEEE Transactions on Control Systems Technology*, vol. 19, no. 3, pp. 707–714, 2011.

[3] C. Huang *et al.*, "Intelligent cruise control of diesel powered vehicles addressing the fuel consumption versus emissions trade-off," in *American Control Conference*, 2018, pp. 840–845.

[4] D. Moser *et al.*, "Flexible spacing adaptive cruise control using stochastic model predictive control," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 1, pp. 114–127, 2018.

[5] P. Golle *et al.*, "On the anonymity of home/work location pairs," in *International Conference on Pervasive Computing*, 2009, pp. 390–397.

[6] J. Williams, "Danger ahead: The government's plan for vehicle-to-vehicle communication threatens privacy, security, and common sense," https://www.eff.org/deeplinks/2017/05/danger-ahead-governments-plan-vehicle-vehicle-communication-threatens-privacy, 2017, accessed: 2017-05-08.

[7] P. Samarati *et al.*, "Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression," in *IEEE Symposium on Research in Security and Privacy*, 1998.

[8] B. Wiedersheim *et al.*, "Privacy in inter-vehicular networks: Why simple pseudonym change is not enough," in *International Conference on Wireless On-demand Network Systems and Services*, 2010, pp. 176–183.

[9] I. Ullah *et al.*, "VBPC: Velocity based pseudonym changing strategy to protect location privacy of vehicles in VANET," in *International Conference on Communication Technologies*, 2017, pp. 132–137.

[10] C. Dwork, "Differential privacy," in *International Conference on Automata, Languages and Programming - Volume Part II*, 2006, pp. 1–12.

[11] Z. Zhou *et al.*, "Differential privacy-guaranteed trajectory community identification over vehicle ad-hoc networks," *Internet Technology Letters*, vol. 1, no. 3, p. e9, 2018.

[12] T. Zhang *et al.*, "Distributed privacy-preserving collaborative intrusion detection systems for VANETs," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 4, no. 1, pp. 148–161, 2018.

Xueru Zhang received her B.Eng. degree in electronic and information engineering from Beihang University, Beijing, China, in 2015. She is currently pursuing her Ph.D. degree in electrical and computer engineering with the University of Michigan. Her research interests include privacy and fairness in machine learning, sequential decision making and distributed optimization.

Chunan Huang received her B.Eng. degree from the Tsinghua University, Beijing, China, in 2016. She is currently pursuing her Ph.D. degree in mechanical engineering with the University of Michigan. Her research interests include modeling, optimal control with applications to transportation systems, including connected automated vehicles and diesel powered vehicles.

Mingyan Liu (M'00, SM'11, F'14) received her Ph.D. degree in electrical engineering from the University of Maryland in 2000. She is current a professor with the Department of Electrical Engineering and Computer Science at the University of Michigan and the Peter and Evelyn Fuss Chair of Electrical and Computer Engineering. Her research interests are in optimal resource allocation, performance modeling, sequential decision and learning theory, game theory and incentive mechanisms, with applications to large-scale networked systems, cybersecurity and cyber risk quantification.

Anna Stefanopoulou received her Ph.D. in Electrical Engineering and Computer Science from the University of Michigan in 1996. She is the William Clay Ford Professor and the Director of the Energy Institute at the University of Michigan. Her research interests include estimation and control of internal combustion engines and electrochemical processes such as fuel cells and batteries.

Tulga Ersal received his Ph.D. in mechanical engineering from the University of Michigan in 2007. He is currently an Associate Research Scientist in the Department of Mechanical Engineering, University of Michigan. His research interests include modeling, simulation, and control of dynamic systems, with applications to transportation and energy systems.
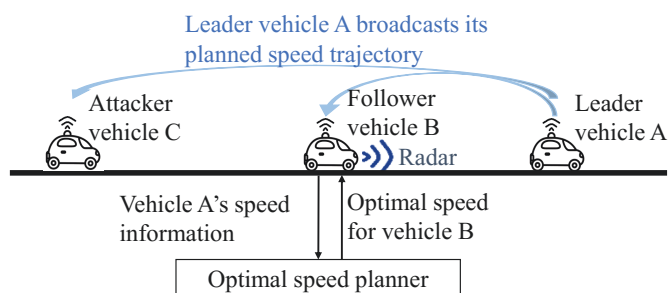
LIST OF FIGURES

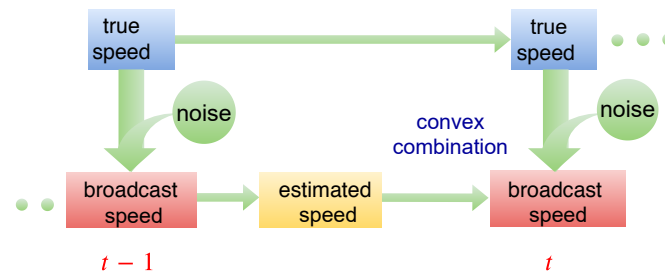Fig. 1. Traffic scenario considered in this paper.

Fig. 2. Two-step illustration of the proposed method: Vehicle A's perturbed speed at each broadcast is determined by the convex combination of the true speed and the estimation from the previous broadcast.
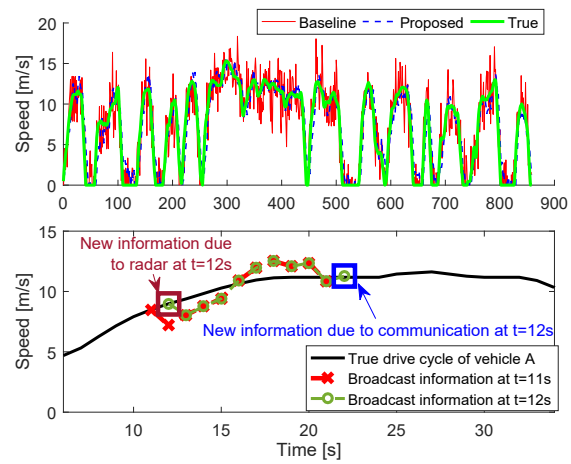
Fig. 3. (Upper) Transmitted information perturbed by the baseline method and the proposed method, and (lower) information available to vehicle B at two consecutive time steps.

Fig. 4. Normalized fuel consumption vs. tailpipe NOx emissions curve for trajectories optimized over the FTP drive cycle with different weight factors w when the optimal speed planner is using speed profiles from vehicle A with different privacy guarantees. The nominal case refers to the scenario when vehicle B follows the same speed trajectory as vehicle A.
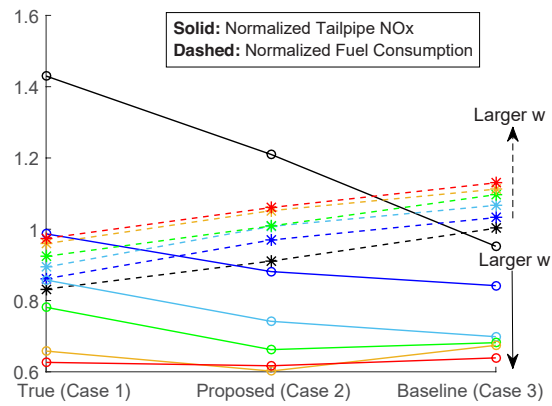
Fig. 5. Variation of normalized fuel consumption and tailpipe NOx emissions with optimal speed planners with different weight factors *w* when information accuracy is changed.