Poster: Distributed and Privacy Preserving Routing of Connected Vehicles to Minimize Congestion

Surabhi Boob*, Shakir Mahmood+, and Muhammad Shahzad*
*Department of Computer Science, North Carolina State University, USA
+National University of Sciences and Technology, Pakistan

***Abstract*—The existing commercial and research based approaches of calculating routes for vehicles suffer from one or more of the following two limitations: they are not privacy preserving and they require expensive infrastructure such as road side units (RSUs). To address these limitations, we propose a distributed and privacy preserving routing protocol, namely DPR, which the connected vehicles collaboratively execute to calculate fast routes to their destinations while preserving vehicle privacy. Our results show that DPR reduces the average travel time of vehicles by up to 42%.**

I. Introduction

Connected vehicles are no longer a futuristic vision. These vehicles can connect to the internet and can also communicate with each other directly. For example, Qualcomm, Audi, and Ford have recently unveiled a communication platform for vehicle to vehicle (V2V) collision avoidance and improved road safety. Similarly, the newer BMW vehicles already directly communicate with each other to automatically prevent collisions in scenarios such as blind intersections. Several standardizations efforts, such as dedicated short range communication (DSRC) [1] are already under way to standardize protocols for direct V2V communication. With such a large number of connected vehicles on the roads in near future, there is an opportunity to leverage their connectivity, and design protocols that these vehicles execute to calculate fast routes to their destinations such that the overall congestion on the given road network is significantly reduced.

II. Limitations of Prior Art

Currently, there are two dominant methods that drivers use to calculate routes to their destinations: cloud based and navigation device based. In the cloud based method, a driver provides his/her destination address to a cloud service, such as Google Maps, using an internet connected device such as a phone and the cloud service sends the current fastest route to the user while taking into account the current traffic conditions. The biggest criticism they face is that they infringe on user privacy by keeping records of the places the users visit [2]. Several surveys have shown that drivers are uncomfortable in providing their destination addresses to these cloud services and would prefer alterative systems where they do not have to disclose their destination addresses to a third party every time they have to travel [3].

In the navigation device based method, the driver enters the destination address into the built-in navigation system of the vehicle or into a navigation gadget, which calculates the optimal route to the destination without sending any information to any third party. Most modern navigation systems have access to the current state of congestion on road networks through the frequency modulated traffic message channel (TMC), which collects latest traffic information through sources such as traffic control centers, traffic cameras, and traffic speed detectors [4]. By leveraging the information from the TMC, these navigation devices calculate fastest routes quite accurately. The key advantage of this method over the cloud based method is that it alleviates the privacy concerns associated with the cloud based method. Unfortunately, the key limitation of this method is that such navigation systems do not calculate routes in a way that would reduce overall congestion on the road network because they do not have information about the routes the other vehicles are taking to reach their respective destinations.

Prior work has also leveraged V2V communication to calculate routes with the objective of minimizing congestion. Existing such schemes have either or both of the following two limitations. First, almost all of the existing schemes assume that dedicated road side units (RSUs) are installed on all the roads [5]. The responsibility of RSUs is almost always to communicate with each other and collaboratively guide vehicles towards their destinations. Unfortunately, such an infrastructure of RSUs incurs very high capital and operational expenses [6], and therefore, has not yet been widely deployed by any city in the world. Second, while several schemes have been proposed that deter the breach of privacy, they do not completely eliminate it because all of these schemes require the vehicle to share its destination address with some other entity such as RSUs, a trusted authority, or other vehicles. As long as a vehicle is sharing its destination address with any other entity, there is a risk of hackers compromising that entity and obtaining information about destinations of the vehicles. The only way to achieve true privacy is to never share the destination address with any other entity outside the vehicle itself. This, however, makes the task of reducing congestion on the road network very challenging.

III. Problem Statement

To address the limitations of the two dominant methods of calculating routes as well as of the prior research, our objective is to design a protocol that leverages V2V communication to calculate fast routes for vehicles to their destinations such that the overall congestion on the given network of roads is significantly reduced, while ensuring that the protocol is 1)

Annex to ISBN 978-3-903176-28-7© 2020 IFIP

distributed, 2) does not require RSUs, and 3) preserves privacy by never requiring any vehicle to share its destination address with any external entity.

Threat Model: The attacker's objective is to determine the destination of any given vehicle. We assume that the attacker can overhear communication between any vehicles in its vicinity. If the communication is encrypted, such as using SSL, the attacker can decrypt it. However, the attacker cannot hack into the internal computing devices of any vehicle. The attacker can also pretend to be a regular vehicle and participate in the V2V communication and calculate a route.

IV. Proposed Approach

In this paper, we present Distributed and Privacy preserving Routing (DPR) protocol that each connected vehicle repeatedly executes in collaboration with the vehicles in its vicinity to calculate fastest route to its destination such that it contributes as little as possible towards causing congestion on the road network. DPR protocol assumes that the vehicle it is executing in has access to the map of the road network, the current GPS location of the vehicle, and the current traffic conditions (obtained via TMC). DPR protocol relies entirely on V2V communication and the already prevalent TMC. It does not require any additional infrastructure such as RSUs.

In any given vehicle that has to travel to a destination, the DPR protocol first locally calculates the fastest route for that vehicle while taking into consideration the current traffic conditions. Next, the vehicle starts traveling towards its destination on this route. As each vehicle has access to the map of the road network and knows its current location, the instance of the DPR protocol running in any given vehicle knows when that vehicle is approaching an intersection. We define an intersection as any place on the road where a vehicle has a choice to take more than one outgoing roads.

Every time the vehicle approaches an intersection, the DPR protocol executing in that vehicle communicates directly with the vehicles in its vicinity, and based on the current traffic conditions, estimates how much those vehicles will contribute to the congestion on the road network if they continue to follow the routes that they are currently following. Using these estimates, the instances of DPR protocol running in these vehicles calculate new fast routes for these vehicles using a distributed algorithm such that the aggregate contribution of these vehicles to the congestion on the road network is minimized. Each vehicle then travels towards the next intersection on its newly calculated route. The motivation behind recalculating the route for any given vehicle *every time* it approaches an intersection is to ensure that the route it follows is optimal under the latest traffic conditions. The motivation behind recalculating the route *only* when the vehicle approaches the intersection is that a vehicle can change its route only from an intersection. Note that the new route that the DPR protocol calculates for any given vehicle can be the same as the route that the vehicle was already following, especially when the congestion on the road network is minimal. In executing the distributed algorithm, the vehicles never share their destination information with each other, which makes the DPR protocol privacy preserving.

The key difference between our DPR protocol and a navigation device based method that can also recalculate route for a vehicle at every intersection is that the navigation device based method is *reactive* while DPR protocol is *proactive*. The navigation device based method reroutes a vehicle only *after* the congestion has happened on the road network. Contrary to that, DPR protocol anticipates congestion and reroutes vehicles *before* congestion even happens, and thus proactively prevents/reduces the congestion on the road network.

V. Evaluation

In this section, we evaluate our DPR protocol using SUMO, a software for simulation of urban mobility [7]. We chose SUMO due to its following three properties. First, it allows us to implement custom routing protocols that the simulated vehicles execute to calculate their routes. Second, it enables us to implement TMC. Third, it allows us to import real world maps from OpenStreetMap [8] and simulates the traversal of vehicles on the roads of the imported map while incorporating the speed limits.

We evaluated DPR under two different traffic intensities, which we will refer to as low traffic intensity and high traffic intensity. We controled the intensity of the traffic on the road network by controlling the average number of vehicles that enter the road network per second. In our implementation, the average number of vehicles that entered the road network per second was 10 for low intensity and 30 for high intensity for the first 40 seconds at the start of the simulation.

We evaluate the performance of the DPR protocol in terms of relative decrease in time (RDT) taken by any given vehicle to go from its origin to destination when using the DPR protocol to calculate its route compared to when using the navigation device based method (with TMC) to calculate its route. Let CT and DT represent the times a given vehicle takes to go from its origin to destination in the presence of other traffic when it uses the device based method and our DPR protocol, respectively. RDT for this vehicle is calculated as RDT = (CT-DT)/CT. We do not compare DPR with any other protocol because when calculating routes, all other protocols benefit from either the exchange of destination information with other entities, or the use of RSUs, or both. This makes the comparison of DPR with any such protocols unfair.

A. Aggregate RDT

To evaluate the aggregate affect of DPR protocol on the travel times of vehicles under each of the two traffic intensities, we ran two sets of simulations on our road network. In the first set, we ran two simulations, one for low traffic intensity and one for high. In each of these two simulations, all vehicles used the conventional navigation device based method along with the current traffic information to calculate their routes locally. From this set of simulations, we obtained the CT values for all vehicles. In the second set, we again ran two simulations, one for low traffic intensity and one for high. In each of these two simulations, all vehicles used DPR protocol to calculate their routes collaboratively. For these two simulations, we ensured that SUMO generated exact same number of vehicles on exact same times with exact same origins and destinations as in the corresponding two simulations in the first set. Thus, from this set of simulations, we obtained the DT value for each vehicle. Consequently, we obtain an RDT value for each vehicle.

Figure 1 shows a box-plot of RDT values for both traffic intensities on the road network. The box-plot corresponding to any given traffic intensity is made using the RDT values of vehicles obtained from the two simulations (one with DPR protocol and the other with the conventional method) under that traffic intensity. Our DPR protocol reduces the average travel times of vehicles by around 25% for high traffic intensity and by over 20% for low traffic intensity. This is a significant improvement considering that it resulted by simply leveraging the V2V communication and without spending multi-million dollars to physically upgrade the road network, such as by increasing the number of lanes on different roads. Adding two lanes to an existing two-lane road costs a whopping \$3.6 million per mile [9]. RDT is higher under high traffic intensity, which shows that DPR protocol is more effective in high congestion scenarios. Figure 2 plots the CDFs of the absolute times the vehicles took to go from their origins to destinations during these two sets of experiments. We calculated from these CDFs that 60% of vehicles experienced an RDT of around 40% for high traffic intensity and 35% of vehicles experienced an RDT of around 55% for low traffic intensity. One might note that the reduction in travel times is in the order of minutes. This is because the distances travelled by the vehicles are not too long, because our aim was to emulate the distances people drive during their commutes in daily life.

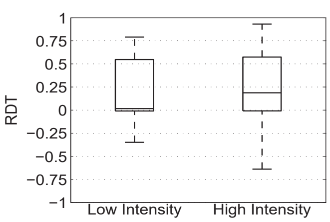

Fig. 1. Aggregate RDTs of vehicles

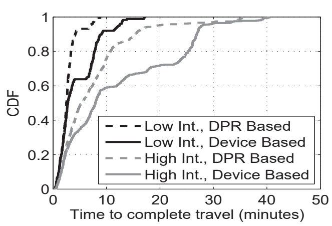

Fig. 2. CDFs of vehicle travel times

B. Effect of Distance on RDT

Next, we show RDT values of vehicles based on how far their destination is from the origin. We observed from our experiments that the closest destination any vehicle had was 250 meters from the origin of that vehicle, and the farthest destination any vehicle had was 4750 meters from the origin of that vehicle. The results we present here are obtained from the same two sets of experiments that we described in Section V-A. Figure 3(a) shows the average RDT values for vehicles that travel different distances on the road network when the traffic intensity is low. Figure 3(b) shows the same when the traffic intensity is high. Any data point corresponding to a distance of x meters, where x is the value written on the x-axes of these figures, is the average of the RDT values of vehicles whose shortest distances to their corresponding destinations lie in the range $[x, x+250)$ in the case of low intensity traffic and $[x, x+750)$ in the case of high intensity traffic.

We observe from these figures that the DPR protocol reduces the travel time of vehicles more under the high traffic intensity compared to under low traffic intensity. This is because their is already less traffic in the low intensity case and the vehicles are not facing much congestion. We also observe that the DPR protocol reduces the travel times of vehicles more when the vehicles travel longer distances. This is because on longer routes, there is higher chance that at some point during its journey, a vehicle will experience congestion, and the DPR protocol will get the opportunity to reduce the congestion.

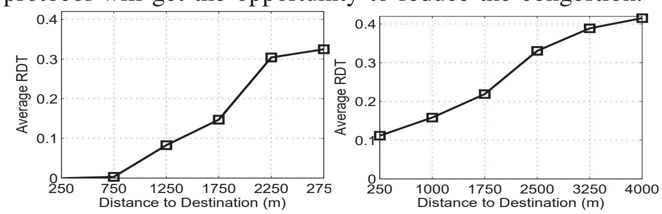

(a) Low intensity (b) High intensity

Fig. 3. Effect of the distance to destination on RDTs

VI. Conclusion

In this paper, we proposed a distributed and privacy preserving routing protocol that each connected vehicle executes to calculate route to its destination in collaboration with several other connected vehicles such that each vehicle contributes as little as possible towards congestion on the road network. The key novelty of DPR is that it does not require any infrastructure support and still reduces the travel times of vehicles significantly while preserving their privacy. The key technical depth of DPR is in the design of the algorithms that the master and participant vehicles use to update routes in a privacy preserving manner. We implemented DPR on SUMO and extensively evaluated it on a real road network under different traffic conditions. We observed from our experiments that the DPR protocol reduces the *average* travel times of vehicles that travel 4000 meters or longer by up to 42%.

Acknowledgment

This work is supported in part by the National Science Foundation under grants # CNS 1616317 and CNS 1616273.

References

[1] J. B. Kenney, "Dedicated short-range communications (dsrc) standards in the united states," *Proceedings of the IEEE*, vol. 99, no. 7, 2011.
[2] "Google Maps Timeline," https://www.google.com/maps/timeline.
[3] "Your car may be invading your privacy," http://www.usatoday.com/story/money/cars/2013/03/24/car-spying-edr-data-privacy/1991751/.
[4] "Traffic Message Channel," http://en.rfwiki.org/wiki/Traffic_message_channel.
[5] J. W. Wedel, B. Schünemann, and I. Radusch, "V2x-based traffic congestion recognition and avoidance," in *Pervasive Systems, Algorithms, and Networks (ISPAN), 2009 10th International Symposium on*. IEEE, 2009.
[6] A. B. Reis, S. Sargento, F. Neves, and O. K. Tonguz, "Deploying roadside units in sparse vehicular networks: what really works and what does not," *IEEE Transactions on Vehicular Technology*, vol. 63, no. 6, 2014.
[7] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz, "Sumo–simulation of urban mobility: an overview," in *Proceedings International Conference on Advances in System Simulation*, 2011.
[8] M. Haklay and P. Weber, "Openstreetmap: User-generated street maps," *IEEE Pervasive Computing*, vol. 7, no. 4, pp. 12–18, 2008.
[9] https://capitolfax.com/summary.pdf.