Mobility21

A USDOT NATIONAL UNIVERSITY TRANSPORTATION CENTER

Carnegie Mellon University

Penn University of Pennsylvania

The Ohio State University

CCAC

INTEGRATION OF AUTONOMOUS VEHICLES WITH ADAPTIVE SIGNAL CONTROL TO ENHANCE MOBILITY

Stephen Smith
Carnegie Mellon University

https://orcid.org/0000-0002-7053-3166

FINAL RESEARCH REPORT

Contract # 69A3551747111

Integration of Automated Vehicle Sensing with Adaptive Signal Control for Enhanced Mobility

Final Report

Stephen F. Smith (PI), Allen Hawkes

Carnegie Mellon University*

Overview

This report summarizes the results of a one-year project aimed at exploiting vehicle-to-infrastructure (V2I) communication to enhance the effectiveness of real-time adaptive traffic signal control systems. As originally formulated, the project's goal was to explore the potential of using the sensing capabilities of connected autonomous vehicles (CAVs) to detect other vehicles in close proximity and use this information to "virtually increase" the level of penetration of connected vehicles in the traffic network, and enhance the predictive accuracy of real-time traffic signal control. However, following initial discussions with project partners Rapid Flow Technologies Inc., provider of the surtrac adaptive traffic signal control system , and Argo AI, an autonomous vehicle technology company, the project focus was shifted to a problem of more immediate and pragmatic importance: a field demonstration and analysis of the ability to further optimize traffic signal control performance through vehicle-to-infrastructure (V2I) communication of real-time CAV route information. In [Hawkes 2016], a mechanism was proposed for incorporating this information into surtrac to reduce uncertainty and generate more accurate predictions of vehicle flows through a controlled traffic network. A benefits analysis of this mechanism, conducted using a microscopic traffic simulation of various traffic networks, showed that network delay was substantially reduced for those vehicles willing to share their routes, and moreover, there was little adverse effect (and even some benefit) to those vehicles not sharing route information.

Building on these ideas, this project has aimed at demonstrating and evaluating V2I route sharing performance in the field and further validating these claims. A cloud-based mechanism for sharing Argo vehicle routes with surtrac in real-time was developed, and appropriate extensions to the current commercial implementation of surtrac were made to factor this additional information into traffic signal control decisions. A pilot test experiment was then designed and carried out using Argo AI test vehicles within the Pittsburgh surtrac deployment. Comparative vehicle delay data were collected along the various routes driven, both with and without route sharing enabled, and results corroborate the benefits predicted by the earlier simulation analysis. Argo vehicles experienced an average reduction in delay of 20% when they shared their routes, and additional analysis of overall surtrac network delay indicated essentially no change in travel time performance to other vehicles. These results are important in that they show that directly benefiting a subset of vehicles does not have to be a zero-sum game, and enable a new, more sustainable model for upgrading urban infrastructure and improving urban mobility through voluntary tolling at the intersection.

* This research was performed in collaboration with Rapid Flow Technologies, Inc. and Argo AI.

Problem

The emergence of technologies for V2I communication offers unprecedented opportunities to improve the performance of real-time traffic signal control systems. The impact on sensing abilities, for example, will be transformational once there are sufficient numbers of connected vehicles (CVs) on the road. Instead of relying on sensing technologies that detect the presence of approaching traffic only at particular locations in the roadway and are prone to significant uncertainty, the traffic control system will have access to second-by-second information of the location, heading and speed of all travelers approaching an intersection. Several recent studies have analyzed the potential benefits of this future state [Beak et.al, 2017, Feng et.al, 2015, Hu, et.al, 2019, Liang et. al., 2019, Tachet et.al 2016]. Unfortunately, this future state is still likely decades away.

This fact has driven our research to investigate questions of (1) whether there are ways to use V2I communication together with adaptive signal control to enhance mobility in the shorter term, when the number of connected travelers on the road is small, and (2) whether such enhancements might in fact serve to accelerate the pace at which CV technology is adopted by travelers. It turns out that the answer to the first question is unquestionably yes. In [Smith 2020], several possibilities are identified, including

- use of real-time bus information to more accurately predict bus arrival times at the intersection and the delays propagated to other following vehicles in the case of near side bus stops, enabling traffic signal control decisions that give priority to transit when appropriate while continuing to optimize other approaching traffic flows (unlike current signal transit priority systems)
- similar use of V2I communication with various municipality vehicles to prioritize their movements in appropriate situations (e.g., giving green to snow plowing vehicles during a snowstorm, expediting emergency vehicles in transit to an event)
- support for safe intersection crossing by pedestrians with disabilities, using smart phone communication to the intersection to signal presence, ensure adequate crossing time, monitor crossing progress, and dynamically extend walk time when appropriate (see [Smith et.al 2019] for further details)
- use of route information communicated by a CV to the traffic signal system to expedite that vehicle through signalized networks

This last possibility involving routing sharing, which was first explored in [Hawkes 2016] and has also been the focus of this project, is unique among the others identified in that it also offers a potential answer to the second question stated above. Imagine that you are a first or last mile freight company. Your delivery routes through the city are known and if your fleet is willing to share them with the traffic control system, overall operations will be improved. Suppose alternatively that you are a high value-of time (VOT) driver. If you are given an opportunity to reduce your travel time by sharing your route, even if you have to pay for it, it may be quite worth your while to do it. Finally, suppose you are the municipality. Your ability to upgrade and maintain your traffic signal infrastructure depends on funding, and in the US this is typically a slow top-down process flowing from the federal gas tax. Each of these separate players stand to gain if the benefits of route sharing were to be made available and brokered via a voluntary “tolling at the intersection model”. Hence there is real incentive for becoming a connected traveler, which in turn can result in faster, broader adoption of CV technologies (with all of their additional benefits).

In this report, we describe our efforts to take the next step toward realizing this voluntary tolling at the intersection concept: that of demonstrating and evaluating the viability of vehicle route sharing in the field. We begin by summarizing the basic route sharing concept first introduced in [Hawkes 2016]. Then, we describe the field implementation that was developed, covering both the cloud-based framework that was used to communicate Argo AI vehicle routes to the surtrac system and the mechanics of injecting route information into the predictive model of approaching traffic constructed by surtrac to provide the basis for generating signal timing plans on each planning cycle. Next, we describe the field test experiment that was carried out and report results that corroborate and validate the earlier simulation analysis. The report concludes by returning to our voluntary tolling at the intersection vision and discussing a path to making it a reality.

Route Sharing Concept

The basic concept of route sharing is to use vehicle route information (e.g. as might be provided by apps such as Google Maps or Waze) to make better real-time traffic signal control decisions. It involves V2I communication of a vehicle’s route through a surtrac -controlled traffic signal network to the signal system, and then subsequent use of this information to improve traffic network performance. The reason to expect performance benefits is straightforward — with the receipt of route information from any given vehicle, uncertainty in the signal system’s predictive model is reduced. For example, the system no longer has to guess whether the vehicle will turn left or go straight at the next intersection. The vehicle has told it which way it will go. Consequently, the traffic signal system can do a better job of optimizing relevant signal timing plans. Because the uncertainty that is reduced centers around those vehicles that share their routes, it would be expected that those vehicles will benefit more than those vehicles that are not sharing. But all vehicles will stand to receive some benefit as the overall performance of the network rises.

A route sharing model was first incorporated into surtrac and evaluated in [Hawkes 2016]. Technically, the approach uses the vehicle route as a means for more accurately projecting the vehicle’s arrival time at downstream intersections, and then updates surtrac’s current predictive model of approaching vehicles to reflect this more accurate information. Since this predictive model of approaching vehicles for any given intersection consists of sequences of vehicle clusters (queues and platoons) at or approaching the intersection from different directions, the adjustment amounts to weighting those clusters that contain each route sharing vehicle at downstream intersections along its itinerary to reflect their greater certainty of arriving at the projected time. Experiments were performed with this route sharing model on a variety simulated traffic networks (using the VISSIM microscopic traffic simulator[1]), and results confirmed expectations. As shown in Figure 1 below (taken from [Hawkes 2016]), a vehicle that shares its route was observed to get through the network substantially faster (on average over 20% as an early adopter), without adversely affecting those vehicles that were not sharing routes. Further, as the percentage of vehicles sharing routes was increased, the performance of the overall network was observed to rise, reducing the delay of all vehicles on the road whether they are sharing routes or not.

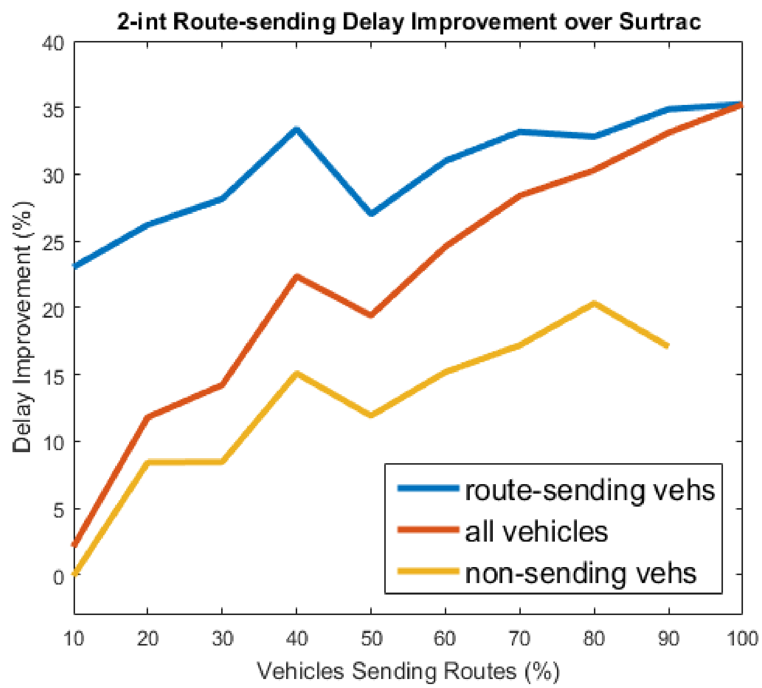

Figure 1: Route sharing results from [Hawkes 2016]

In producing these results, experiments were performed with differing amount of forward route projection (i.e., how many intersections forward to do cluster adjusting from the vehicle's current location). It was found that factoring in more than 2 downstream intersections yielded negligible further improvement, most likely due to the fact that that placed the vehicle outside of surtrac's current prediction horizon.

Field Implementation

To develop a field implementation for route sharing, we start with the assumption that participating vehicles will interact with a cloud-based server process. Although direct

[1] https://www.ptvgroup.com/en/solutions/products/ptv-vissim/

communication with individual intersections could be a feasible alternative, we believe a cloud-based server process is more logical and efficient. In normal operations, we envision that these vehicles would initiate a route sharing request by sending a request message to the infrastructure's cloud-based server. The vehicle would have a communication device provided for this purpose (cellular V2I), and the vehicle would know the infrastructure server address as a result of signing up for the route sharing service. The request message would transmit the route as a sequence of way points, and also indicate the vehicle's current location.

In this initial pilot experiment, however, we were able to take advantage of the fact that Argo AI already has cloud-based servers to track its vehicles. For each Argo AI vehicle that is currently on the road, a representation of its route and current location is continuously maintained and updated. Given this pre-existing infrastructure, we developed a cloud-based, route sharing server interface that periodically polls the Argo servers, and then carries out the envisioned route sharing API on behalf of each Argo AI vehicle. For this initial field test, we explore the benefits of propagating route information just one intersection ahead.

With this structure in place, the following concept of operations was implemented:

- Step 1: A new route sharing request is received by the cloud server, consisting of a vehicle message specifying a sequence of future waypoints and the vehicle's current location.
- Step 2: The route sharing cloud server process extracts the sequence of way points and the vehicle's current location from vehicle message and maps waypoints to a set of surtrac -enabled intersections.
- Step 3: The server process then computes the vehicle's expected arrival time at the next intersection along its route, using a free flow time calculation based on the expected speed of the vehicle if there are no blocking vehicles. This expected arrival time is then communicated to the surtrac process running at that intersection, along with a request to provide an expected departure time for the vehicle at that intersection.
- Step 4: the surtrac process running at the next intersection intersection updates its predictive model of approaching vehicle clusters to include a new cluster representing the route sharing vehicle, and associates a weight with this cluster that is reflective of the greater certainty that is associated with this vehicle's arrival. Subsequent online planning cycles at this intersection will generate signal timing plans that incorporate this new knowledge. The route sharing vehicle's departure time will then be communicated back to the route sharing cloud server.
- Step 5: When it is recognized that the vehicle has traversed the next intersection, The route sharing server process proceeds to check if the intersection just traversed was the last surtrac -enabled intersection in the vehicle's route. If it is, vehicle delay metrics are captured and the process is terminated. Otherwise, the server re-initiates the sequence of steps starting with Step 3 above, updating the predicted arrival time of the vehicle at the new next intersection.

Results

To evaluate this field implementation and quantify the benefits of route sharing in the field, an on-off pilot study of the implemented capability was conducted within the Pittsburgh surtrac deployment. The study was conducted over the period from February 11 to February 21, 2020. The analysis was restricted to weekdays during this period, and Argo AI vehicles moving through the surtrac deployment were constrained to operate with route sharing enabled one day, then disabled the next and so on. Participating Argo AI vehicles simply proceeded as normal with their daily road testing in the East End, which included 13 surtrac -controlled intersections that were enabled for route sharing. Over the course of the 2-week pilot testing period, the routes of 15 distinct Argo AI vehicles were utilized, although not necessarily at the same time. Across all runs, vehicle travel time information was collected, both for the Argo AI vehicles serving as test vehicles individually, and for the overall surtrac controlled traffic network. Overall, data was collected for 5 days with route sharing on, and 4 days with route sharing off.

Examination of the vehicle trajectories produced during the evaluation period revealed a number of routes where Argo vehicles unexpectedly stopped and failed to advance for an extended period of time, despite not being blocked by traffic signals. These stops were confirmed to be caused by circumstances unrelated to the experimental study. Given this understanding, such anomalous route segments were filtered out of the collected data before computing performance results. Considering this filtered data set, the use of route sharing was found to result in an average reduction in delay of 20.3%, and a 14.7% reduction in trip travel time. The details are shown in Table 1, where delay is defined as: *delay = travel_time – freeflow_time*. These results largely corroborate those obtained originally in simulation at low levels of CV penetration.

Mode	Number of Intersections traversed	Distance per intersection (meters)	Delay per intersection (seconds)	Travel time per intersection (seconds)
Route sharing ON	76	93.8	18.7	27.2
Route sharing OFF	63	92.7	23.4	31.9

Table 1: Comparative results, with and without route sharing.

To confirm that these results were not unduly influenced by different traffic volumes, a time of day comparison of intersection traversals was also tabulated, which is depicted in Table 2. These results include all intersection traversals, regardless of whether the trip to which it belonged experienced anomalous stops. This data indicates a reasonable balance in trials during the two highest volume hours of the daily testing period from 9-11am.

To get an indication of the impact of route sharing on the traffic flow performance of the surrounding, non-participating vehicles, we also examine the delay per vehicle network wide. To do so, we make a few assumptions. First, we restrict our analysis to the busier morning travel period from 8AM to 11AM where potential negative effects on overall traffic flows is likely to be the worst. Second, we include vehicle delay information from the full surtrac network of 13 intersections that was utilized in the evaluation, to account for possible indirect network effects (e.g., delays at downstream intersections). Third, we use vehicle delay data collected from 3PM – PM, the non-testing portion of each testing day, to establish a baseline performance difference for comparing days, apart from the introduction of route sharing.

Mode	**8am**	**9am**	**10am**	**11am**	**12pm**	**1pm**	**2pm**
Route sharing ON	32	22	32	0	27	25	2
Route sharing OFF	48	25	29	0	2	12	0

Table 2: Time of day comparison, with and without route sharing.

Table 3 below shows the difference in average network delay between days where route sharing was enabled (Jan 11, 13, 17, and 19) and those where it was not (Jan 12, 14, and 18) over the morning portion of the testing interval (8AM – 11AM). Table 4 shows the baseline difference in network level delay for the same sets of days, using the non-testing portion of each day. Overall, we see that route sharing caused a 2.8% increase in network wide delay while achieving a 20.3% reduction in delay for those vehicles sharing their routes. However, analysis of the non-testing portion of the day for the same sets of route sharing on and off days, we see that over the actual days tested, the "route sharing off" days had a 9.7% better baseline (i.e., surtrac only) network performance than that of the "route sharing on" days. Hence, the negative impact of route sharing on overall network performance appears to be negligible.

Time Range	**Dates Analyzed**	**Control**	**Number of Vehicles**	**Total Delay (Mins)**	**Delay Per Vehicle (Secs)**
8AM – 11AM	Jan 11,13,17,19	surtrac w/ route sharing	4258	1940.6	27.3
8AM – 11AM	Jan 12,14,18	surtrac	2991	1266.3	26.6

Table 3: Average delay performance of full network with and without route sharing enabled

Conclusions and Discussion

The goal of this project has been to verify that the benefits of V2I route sharing in conjunction with real-time schedule-driven traffic control. Working in collaboration with Rapid Flow Technologies and Argo AI, an initial field experiment was designed and carried out, utilizing the Pittsburgh surtrac controlled traffic network, and a cloud-based infrastructure for route sharing that was implemented to provide V2I linkage of Argo AI vehicles to surtrac traffic control decisions. The field experiment produced results comparable to those reported in prior simulation studies, and has clearly demonstrated the potential of V2I route sharing. On the basis of these results, Rapid Flow Technologies has subsequently carried out a more structured evaluation using municipal vehicles in the Quincy MA surtrac deployment. In this case, the same sets of routes were driven with and without the use of route sharing by local municipality workers equipped with an experimental smart phone app to communicate to the cloud server. This study showed a 35.4% average reduction in delay to vehicles willing to share their routes.

Time Range	Dates Analyzed	Control	Number of Vehicles	Total Delay (Mins)	Delay Per Vehicle (Secs)
3PM – 8PM	Jan 11,13,17,19	surtrac	11396	6622.9	34.9
3PM – 8PM	Jan 12,14,18	surtrac	8185	4335.2	31.8

Table 4: Average delay performance baseline of full network

Given the effectiveness of routing sharing in the field, the logical next question to ask is how best to deploy and utilize this capability. We believe that route sharing offers an unprecedented opportunity to accelerate deployment of adaptive signal control technology in cities, vastly improving urban mobility in shorter time frames than is possible today, while at the same time providing a more sustainable approach to managing transportation infrastructure. Specifically, we envision an operational model that incorporates *voluntary tolling* at the intersection, utilizing the travel efficiency boost provided by route sharing as incentive for participating. Increasing cities are resorting to tolling as a means of curbing congestion, but these schemes really take a “blunt hammer” approach, uniformly taxing all vehicles on the road for the privilege of traveling into the city and denying travelers based on their ability to pay. The concept of voluntary tolling, in contrast, is based on incentivizing travelers to pay, and offering value in return for participation. A vehicle pays a few pennies each time it traverses a signalized intersection in return for the ability to share its route with the traffic signal system, which in turn provides value to the vehicle in terms of shorter travel times. Imagine a first/last mile freight company that is operating in a city, possibly with a distribution center nearby. The company knows the routes that its fleet of vehicles will drive on any given day, and if the company opts in, it can immediately accrue cost savings greater than the tolling fees it expends to participate through improved operational efficiency. There are similar incentives for high value-of-time travelers to participate, as well as city transit services and other municipality fleets (e.g., snow plow vehicles, emergency vehicles). And remember, as more and more travelers opt in to voluntary tolling, analysis shows that the overall performance of the traffic signal control network will continue to rise, reaching a performance ceiling only as the level of participation approaches 100%. So all vehicles on the road receive benefit, with participants obviously receiving the lion's share.

From the municipality's perspective there is a different benefit. Our vision of voluntary tolling at the intersection turns the traditional infrastructure funding model on its head, and instead promotes a more agile, bottom-up process that expedites introduction of new technology and creates a much more sustainable approach to transportation infrastructure. The prospect of voluntary tolling enables new infrastructure acquisition possibilities. For example, a municipality might procure adaptive traffic signal control technology city-wide under terms that payment could be made on the backend, from revenues accumulated by the voluntary tolling model. Once the procured technology has been acquired, the municipality inherits a continuing revenue stream that can be used to maintain and upgrade transportation infrastructure over time.

Congestion in urban environments grows costlier every year and with continuing trends toward urbanization throughout the world, the need for new perspectives and new approaches to urban mobility is crucial. In that regard, we believe that the concept of voluntary tolling at the intersection, incentivized by the boost in traffic flow efficiency that route sharing can enable when coupled with real-time adaptive signal control, provides an unprecedented opportunity to revolutionize urban traffic flow. First, it will reduce travel delays through signalized networks for all travelers by 30-35% over the conventional fixed signal timing approach to traffic signal control that dominates the transportation infrastructure landscape in cities today[2]; and those travelers that choose to opt into voluntary tolling and share their routes will immediately receive an additional 20-25% reduction in delay. As the level of participation in voluntary tolling rises over time, so will the overall performance of the traffic signal network, providing further efficiency benefits to all travelers. Second, voluntary tolling enables technology acquisition models that promote widespread upfront installation of adaptive signal control technology throughout a city and recovery of costs on the backend with tolling revenues, providing a path to both rapid adoption of advanced traffic signal control technology and a more sustainable, bottom-up approach to infrastructure maintenance. This project has taken the next step toward realization of this concept, confirming that the travel time efficiency improvement that can be achieved by V2I route sharing in the field is substantial, and consistent with what was predicted in earlier microscopic traffic simulation studies.

References

[2] Based on Rapid Flow Technologies' experience in its surtrac deployments in 9 North American cities.

[Beak et.al, 2017] Beak, B., Head, K. L., & Feng, Y., “Adaptive Coordination Based on Connected Vehicle Technology”, *Transportation Research Record: Journal of the Transportation Research Board*, 2017.

[Feng et.al 2015] Feng, Y., Head, K. L., Khoshmagham, S., & Zamanipour, M., “A Real-Time Adaptive Signal Control In A Connected Vehicle Environment”, *Transportation Research Part C: Emerging Technologies*, 55, 460-473.

[Hawkes 2016] Hawkes, A., “Traffic Control with Connected Vehicle routes in surtrac”, Unpublished MS Thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA., 2016

[Hu, et.al, 2019] Hu, H-C, S.F. Smith and R. Goldstein, “Cooperative Schedule-Driven Intersection Control with Connected and Autonomous Vehicles”, *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2019)*, Macau, China, November, 2019.

[Liang et. al. 2019] Liang, X., Guler, S.I. and Gayah, V.V., “Joint optimization of signal phasing and timing and vehicle speed guidance in a connected and autonomous vehicle environment”, *Transportation Research Record: Journal of the Transportation Research Board*, 2673(4):70-83, 2019.

[Smith et.al., 2019] Smith, S.F., Z.B. Rubinstein, J. Marousek, M.B. Dias, and H. Radwanick. “Connecting Pedestrians with Disabilities to Adaptive Signal Control for Safe Intersection Crossing and Enhanced Mobility: Final Report”, Federal Highway Administration Joint Program Office Publication FHWA-JPO-19-754, September, 2019.

[Smith 2020] Smith, S.F. “Smart Infrastructure for Future Urban Mobility”, *AI Magazine*, 41 (1), Spring 2020.

[Tachet et.al 2016] Tachet, R., Santi, P., Sobolevsky. S., Reyes-Castro. L.I., Frazzoli, E., Helbing, D., Ratti, C., “Revisiting Street Intersections Using Slot-Based Systems”, PLoS ONE 11(3): e0149607, doi:10.1371/journal.pone.0149607, 2016.