Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnline CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 55-41 (2022) 83–88

Medium-scale to large-scale implementation of cyber-physical human experiments in live trafic.

Sean T. McQuade[⋆], Chris Denaro[⋆], Malaika Mahmood[†], Jonathan W. Lee[§], Gracie Gumm[‡], Jonathan M. Sprinkle[‡], Daniel B. Work[‡], Benedetto Piccoli[⋆¶], Benjamin Seibold[‖], Alexandre M. Bayen[§]

Abstract: Autonomous Vehicles (AVs) such as cars and trucks are being developed and tested as Cyber-Physical Human systems while the technology improves. Before these systems can achieve full autonomy, some serve as tools in the form of adaptive cruise control.

The CIRCLES Consortium investigates the potential for AVs to increase fuel eficiency of highway trafic by smoothing “stop-and-go” trafic waves that result from normal human driving behavior in congestion. We have performed an experiment to evaluate the real world effects of implementing this strategy. A medium-scale experiment was performed on I-24 near Nashville, TN in August 2021. This was a precursor to a larger experiment that will take place November 2022. We examine how the human part of the experiment will change as we scale up from an 11 vehicle test (four AVs) to 100 AVs. There are many solutions to problems of the medium-scale experiment that would be inconvenient, complicate the experience, or not be practicable.

The medium-scale experiment involved 11 cars of which four had a custom control algorithm installed to be engaged by the driver. The large-scale experiment will have 100 cars, all with custom control algorithm installed to act on trafic when the controller is engaged.

We examine key choices made for the medium experiment, and how some will be different for the large experiment. Our experience performing the medium-scale experiment has made it clear that repeating our methods from this smaller one are ineficient or impossible if used for the large-scale experiment and will be improved.

Keywords: cyber-physical human systems, automotive cooperated control (ADAS, etc), trafic waves, autonomous vehicles, trafic control

1. INTRODUCTION

Smoothing stop-and-go trafic waves to increase the energy eficiency of surrounding vehicles has been demonstrated in several studies, Stern et al. (2018); Jin et al. (2018); Ma et al. (2016). However these have been in simulation or with actual vehicles in simplified situations like a single-lane ring road. There have been large-scale efforts for reconstructing trafic data using a sample of mobile phone GPS in trafic, Herrera et al. (2010), but these efforts sought to describe trafic instead of affecting it.

The next step is to bring these strategies to the actual highway in live trafic. In 2021, the CIRCLES Consortium[1], Bayen et al. (2020, 2021); Sprinkle et al. (2022), performed a medium-scale field test in live trafic on I-24 near Nashville, TN to apply trafic smoothing techniques to a real world situation. This was a precursor to a large-scale experiment to take place later this year. These researchers have done large-scale vehicular data collection and will draw upon this experience such as the Mobile Century project (see Work et al. (2009)) to plan the larger field test. Both the 2021 and 2022 tests will be co-located with the I-24 Motion project, where high resolution cameras are mounted on 110 foot poles every 0.1 miles between Exit 62 and Exit 59, Gloudemans et al. (2020). When fully erected in Q4 2022, the cameras will allow continuous coverage of this part of I-24, capturing approximately 150,000 vehicle trajectories daily.

When fully autonomous vehicles (AVs) become a reality, interspersed AVs could improve the energy efficiency of all the surrounding trafic. Presently, we must test these ideas with humans activating modified adaptive cruise controls (ACC) that attempt to dissipate trafic waves. These tests are performed as a Cyber-Physical Human system given the absence of full autonomy.

[⋆] Sponsor and financial support acknowledgment goes here. Paper titles should be written in uppercase and lowercase letters, not all uppercase.

[⋆]–Center for Computational and Integrative Biology, Rutgers University-Camden, Camden, NJ, USA

[†]–Department of Biology, Rutgers University-Camden, Camden, NJ, USA

[‡]–Vanderbilt University, Nashville, TN, USA

[§]–UC Berkeley, Electrical Engineering and Computer Science, Berkeley, CA, USA

[¶]–Department of Mathematical Sciences Rutgers University-Camden, Camden, NJ, USA

[‖]–Temple University, Philadelphia, PA, USA

[1] https://circles-consortium.github.io/, Nguyen and Lee (2020); circles.consortium@gmail.com

10.1016/j.ifacol.2023.01.107

This research has used the combined skills of many scientists. Mathematicians, engineers, and computer scientists are required to design and implement the control algorithms that may smooth trafic in certain conditions. Mathematical modelling is used to develop fuel rate models of both specific and generic vehicle types. Such models are used in optimizing automated driving behavior. Software engineers are needed to interface with the CAN bus of a vehicle to be able to read measurements and implement control algorithms which give commands to the car. Hardware engineers are required to install and fix components in the cars. Other scientists must also plan and manage these teams to coordinate our efforts. A sizable part of this planning and management is to anticipate the needs and actions of human AV operators as well as the surrounding drivers in trafic.

The AV operators are a key human element in the experiment, and solutions to the medium-scale test may be intractable for the large-scale experiment. A simple example of the dramatic difference between the two experiments is having the operators get in their cars and exit the parking lot with minimal time between each car. With 11 driver-operators this was trivial. To facilitate our platoon entering the highway, we were aided by the Smyrna Police Department. They stopped trafic at a stop sign and two trafic lights, and led the platoon onto the on ramp.

In the large-scale test, 100 AVs should be dispersed among trafic. Even if practicable, a contiguous platoon of 100 cars would not be a desirable experiment; instead we would want them distributed naturally throughout the testbed. Therefore, we will steadily release the AVs from our initial lot, which is expected to take more than 30 minutes.

The rest of this paper is organized as follows: the next section presents the details of the medium-scale experiment, including a discussion of the resulting data analysis. Section 3 establishes the key new challenges to executing the large-scale experiment. Section 4 concludes with the differing approach to driver instruction and training that is necessitated by the difference in scale.

2. THE MEDIUM-SCALE EXPERIMENT

A medium-scale experiment took place from July 28, 2021 to August 6, 2021 (one week of preparation and driver training, and one week of testing). 11 cars were run on I-24 near Nashville between Exit 66 and Exit 57. This experiment served to test the technical and logistical feasibility and to prepare for the large-scale experiment planned for early November 2022. Due to the smaller scale of the 2021 test, drivers were selected primarily from researchers already involved in the CIRCLES project. The four AV drivers were each Ph.D. students who had experience in designing the control algorithms. These drivers hailed from Vanderbilt University, The University of Arizona, and UC Berkeley.

The purpose of the medium-scale experiment was to give us experience in three main areas: (1) to implement control algorithms on ACC systems by sending messages through the CAN bus system of the cars; (2) to measure the effect of this in live trafic to compare energy results to previous experiments, such as tests performed in special situations, e.g. the ring road experiments, Chou et al. (2022); (3) to gain experience in managing the many moving parts of such an experiment in live trafic.

The seven drivers in HDVs were all Vanderbilt University afiliates, ranging in experience from undergraduate students to research scientists. All but one of these students had been involved in the project in technical aspects. The driver that had not previously been involved is an undergraduate civil engineering student at Vanderbilt University and was selected based on her enthusiasm to assess research as a career opportunity. The demographics of the drivers were: 45% of the drivers represented racial minority groups and 45% of the drivers were female. In addition to the drivers, each vehicle contained a passenger, all of whom had been involved in technical aspects of the project. The passengers were stationed in vehicles with a handheld transceiver that maintained communications with the full vehicle fleet. Passengers in the four AVs were selected based on their experience with the control algorithms, the majority of whom had assisted in programming or installing them. Passengers were responsible for relaying information regarding roadway hazards, their location along the testbed, and their return to base. Once all vehicles had stated their return to their designated parking spot, the test was declared finished and occupants were instructed to exit their vehicles. Safety was the top priority in the experiment and the passengers played an integral role in ensuring zero incidents occurred.

The route, shown in Fig. 1, was driven with all vehicles in lane 2 (i.e., the second lane from the left) for the Westbound portion. Four vehicles were equipped with specially designed longitudinal controllers, and the remaining seven were human-driven vehicles (briefly HDVs). All drivers were instructed to get to their target lane safely, and AVs would activate controllers at their operators' discretion after reaching the target lane. All the vehicles exited at Exit 57 and drove Eastbound in any lane that was convenient (the AVs were not running controllers on the Eastbound side of the route). The fleet would exit at off-ramp 66, then regroup at a nearby staging area. The experiment continued to iterate in this pattern until it was clear that the rush hour congestion was clearing up, resulting in three to four drives per testing day.

The AV controllers consisted of a feedforward neural network optimized using deep reinforcement learning algorithms. A simulation was built using real-world trajectory data in which agents were trained in a single-agent fashion to minimize fuel consumption, large gaps, and accelerations amplitude. Fail-safes and hard-coded gap-closing were added atop the model to prevent unwanted out-of-distribution behaviors; more details can be found at Lichtlé et al. (2022). Controllers are tested on an end-to-end integrated simulation pipeline before being implemented in vehicle systems, Lee et al. (2021).

When translating results of learning-enabled controllers that were trained on simulation engines, supervisory and other controllers were deployed for additional safety. Controllers were imported into a Robot Operating System implementation that utilized real-time data acquisition and control, Elmadani et al. (2021), alongside model-based code-generated software, to execute the controller

Fig. 1. The driving route for the 2021 experiment (medium-scale) on I-24 Southeast of Nashville. All 11 cars drove in lane 2 during the Westbound direction. AVs (red) were between manually driven cars (cyan) that were instrumented to measure the AV's effect on trafic.

in real time on the vehicles. Supervisory safety controllers extended from those used in Bhadani et al. (2018) were used to filter control commands prior to transmission to the vehicle.

Each vehicle had a unique designation and the group of 11 vehicles was ordered as follows: HDV1, HDV2, AV3, HDV4, AV5, HDV6, AV7, HDV8, AV9, HDV10, HDV11. Every car had a driver and a co-pilot, which allowed co-pilots to communicate via HAM radios. The range of the radios was suficient for the first car to talk to the last car during the Westbound part of the experiment (the turnaround point at Exit 57 would separate the vehicles, separating some vehicles from radio range). The front car, HDV1, had an experienced driver from Vanderbilt University who knew the route well. The last car, HDV11, had Professor Dan Work as a co-pilot, who led many aspects of the test, including the safety team.

2.1 Data Collection and Analysis

All vehicles are equipped with a comma.ai panda black, a Raspberry Pi, and a personal WiFi hotspot so that GPS and CAN bus data can be stored and uploaded for remote preservation. Data were collected via the Libpanda tool, Bunting et al. (2021), providing a sampling of trafic sensing that can be analyzed *post hoc* as a proof of concept prior to the large-scale experiment.

The first step of the data analysis is to clean the data, which involves the following:

(1) Map matching snaps raw GPS coordinates onto the known street network (obtained from Open-StreetMaps). Only points that are on the actual west-bound portion of the highway and detected to be in the designated lane are kept. The result is a time series of one-dimensional position, where zero denotes the end of the Exit 66 on-ramp.
(2) CAN decoding is performed using the Strym software library, Bhadani et al. (2022), with the team's custom-compiled DBC files, specifically for the Toyota RAV4 models that were used. This step extracts speed, acceleration, space gap, leader relative speed, and ACC engagement status from the CAN bus data.
(3) Outlier removal is performed by comparing GPS data to CAN data. Both GPS and CAN data can contain outliers, or in extreme cases, entire data streams contain non-updating constant values. Data that is incongruous is discarded, to be filled by the following interpolation step.
(4) Signal interpolation serves multiple purposes, namely, (a) filling gaps of missing and/or removed data, (b) enforcing time series alignment onto a fixed time grid of 10Hz since all data streams have different publish frequencies, and (c) projecting CAN system time (i.e., variable offsets) onto the more precise GPS/satellite time.
(5) Signal fusion & smoothing combines the GPS and CAN data for position, velocity, and acceleration. The formulation solves for a ground truth position vector by minimizing the linear combination of position, velocity, and acceleration errors, as well as a term for jerk smoothing. The same procedure is performed for the inferred leader of each vehicle, however in that case, a distinct data source for acceleration is not available.

For the remainder of this section, we present data from one particular drive. In Fig. 2, the space gap (top) and time gap (bottom)[2] of HDVs (cyan) and AVs (red) are displayed in histograms. The space gap distribution shows clearly that the control algorithm tends to keep larger gaps than the HDV drivers. More strikingly, the time gap distribution shows that the AV control policy attempts to maintain tight control around 2.7–3.0 seconds of time gap, contrasted to the broad distribution of time gaps for the HDVs.

In Figs. 3 & 4, we see the trajectories of two AVs plotted side-by-side. The background color, which denotes ACC engagement (green) or disengagement (red), shows that the AV operators disengaged a handful of times each, but was predominantly engaged during this drive. Again, the control policy's tendency to maintain a time gap just below 3.0 seconds is clear in the bottom plot.

3. NEW CHALLENGES AT SCALE

By contrast the large-scale experiment will include 100 AVs, all capable of running controllers. We will rely on the sophisticated camera system instead of HDVs to measure trafic impact. The test will take place across lanes 2, 3, and 4 (i.e., the three right-most lanes, discounting merging lanes). In this experiment, we will do approximately seven laps of the two routes shown in Fig. 5. As in the medium test, the AVs will only have controllers engaged on the Westbound part of I-24. The medium experiment was necessary to understand the planning and logistical hurdles that would accompany the large-scale experiment. The large-scale experiment is designed to control trafic waves across three different lanes, see inset of Fig. 5. This requires us to design the experiment to maximize the penetration rate that allows dampening stop-and-go waves. The penetration rate is the amount of AVs as a percentage of the

[2] Space gap is defined as the distance from the ego vehicle's front bumper to the leader vehicle's rear bumper. Time gap is defined as the space gap normalized by the ego vehicle speed, i.e., the time it takes to cover the distance of the space gap.

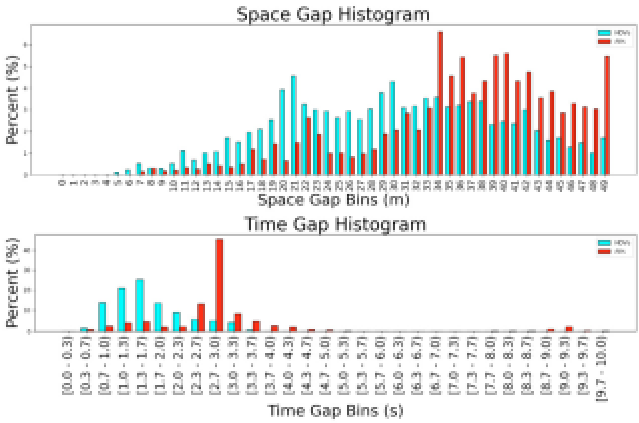

Fig. 2. Histograms of space gap (top) and time gap (bottom) measurements collected from the vehicles. Measurements are aggregated according to vehicle type, i.e., HDVs in cyan and AVs in red. The AV measurements are shown regardless of control engagement or disengagement. Clearly, the AVs tend to keep larger gaps than their HDV counterparts, in particular at the approximate 3-second time gap.

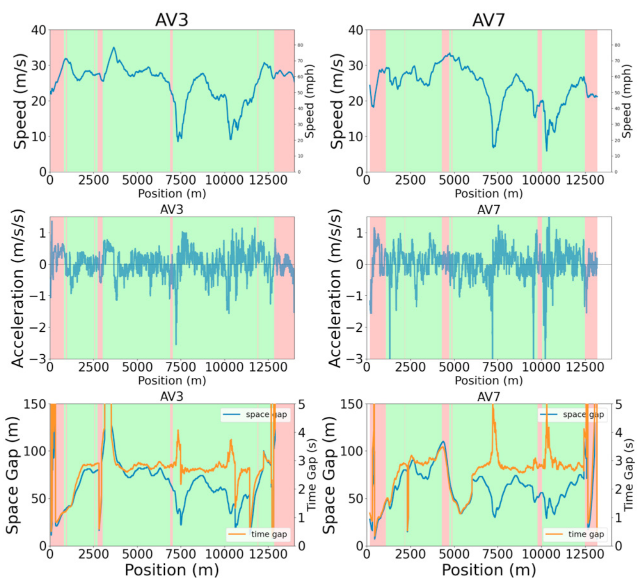

Fig. 3. Example time-series for various metrics of two AVs (AV3 on left and AV7 on right) during one of the tests. The background color denotes whether control was on (green) or off (red), as AV operators were instructed to disengage under unsafe conditions. The rows show speed (top), acceleration (middle), and gap (bottom). The gap plot shows space gap (blue) on the left y-axis and time gap (orange) on the right y-axis. Notably, the control algorithm can be seen to approximately keep a 3-second time gap.

total trafic on the road. In highly specialized examples, such as 22 cars on a single lane ring road, it is shown that a penetration rate of 4.8% can smooth waves, Stern et al. (2018).

Planning vs execution of the medium-scale experiment allowed us to learn what needed improvement and verify various pipelines, such as controller verification and validation testing. We learned that some processes are likely to scale poorly, such as our pre-experiment safety briefing, hardware checks before starting, breaks for drivers in the middle of a long test, fueling the cars, and returning the cars to storage at Vanderbilt after testing. There are also bottlenecks that change the intended test route. An important one is an obvious difference between the medium experiment route shown in Fig. 1 and the large test route shown in Fig. 5. In the 2021 experiment, all vehicles would take the same route: Exit 57 to Exit 66. Exit 57 is an off-ramp that queues to a trafic light. The cars then make a left turn and cross beneath the highway through an underpass to a second trafic light. At this light an unprotected left is made to access the on-ramp for I-24 Eastbound. We realized that scaling to 100 cars, we would significantly extend the queue for the trafic lights.

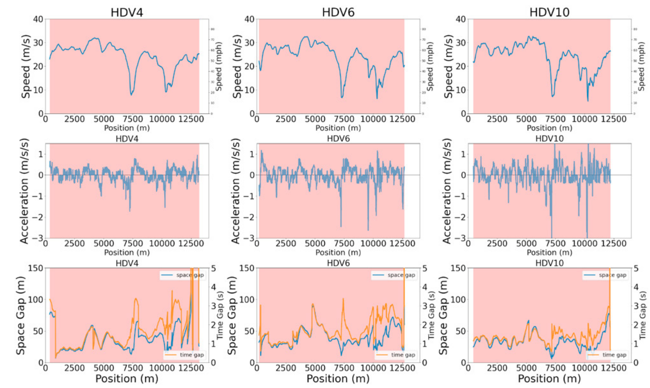

Fig. 4. The same plots as shown in Fig. 3 except for three of the HDVs from the same drive. By contrast to the AVs, humans keep smaller and significantly more erratic gaps.

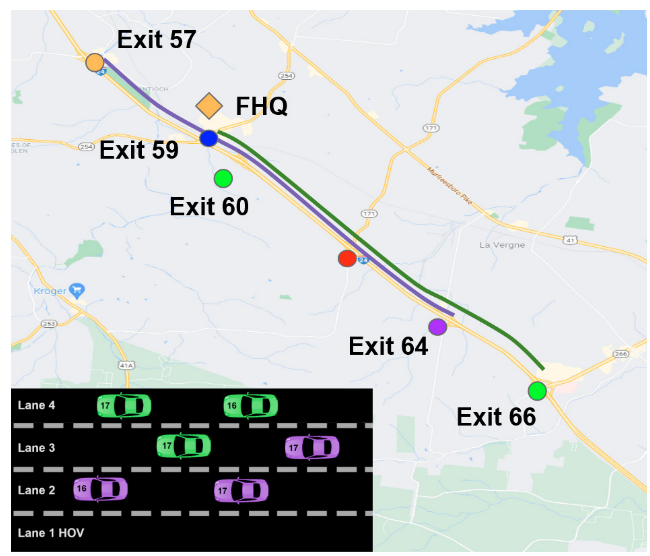

Fig. 5. The two driving routes for the 2022 experiment (large-scale) on I-24 Southeast of Nashville. All cars are AVs and while traveling Westbound have designated lanes 2, 3, and 4. Drivers on both routes would take a scheduled break for a driver switch (or as needed) at the Field Headquarters (FHQ) depicted as an orange square North of Exit 59.

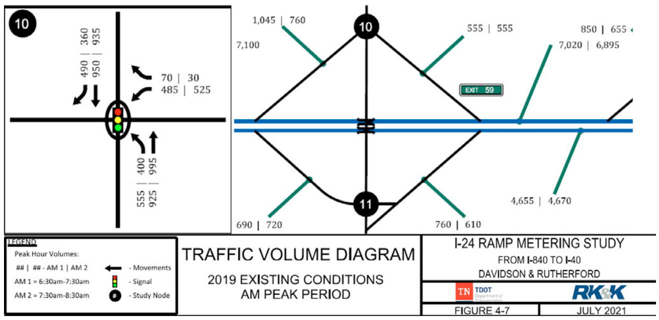

Fig. 6. I-24: On each segment, ramp, and trafic light, two numbers are given. These are the number of cars that passed through the road from 6:30am to 7:30am (left) and 7:30am to 8:30am (right) from a 2019 study.

Using I-24 ramp and mainline trafic volume data during peak congestion–provided by our partners at TDOT–(shown in Fig. 6), we were able to estimate our hypothetical impact on trafic by adding an additional 100 cars during peak hours of trafic. We paid special attention to estimating how we may adversely affect the flow rate at the trafic lights. We combined this flow rate data with trafic light pattern that was active throughout different hours of the day (also provided by TDOT) to estimate the additional trafic that could be added without causing queuing on the I-24 off ramps. We estimated a theoretical maximum number of cars through the light given the amount of green light time compared to our 2021 drives through the same lights. We then estimated the total trafic on those lights by summing the contribution from background trafic and from our 100 AVs.

We calculated that the desired route would take approximately 42 minutes during maximum congestion; this estimate is being checked by real drives through the exact route during peak congestion. With this estimate, AVs would complete 1.43 loops of the route per hour. Therefore, our 100 AVs would really add 143 vehicles per hour to the off-ramp queue. We decided that this was too large compared to the hourly data showing that 355 cars would normally make a left through the first light; 380 cars would normally make a left through the second light. We were concerned that our AVs would cause anomalously long queues to form on the off-ramp at Exit 57 and Exit 59.

To mitigate trafic disruption through the trafic lights we decided to break the route up into two so that only half the AVs would use this exit to turn around. The other half would use the nearest acceptable exit (Exit 59). This is the reasoning behind the two routes shown in Fig. 5. 50 cars will run on the Green route (33 in lane 4, and 17 in lane 3). 50 cars will run on the Purple route (33 cars in lane 2, 17 cars in lane 3). The overlapping section of the two routes corresponds to the location of the new trafic cameras that will capture the experiment.

We will need to practice the briefing and car release from the parking lot before the test. For the medium-scale experiment we gave a safety talk to all drivers and co-pilots before walking to the vehicles. Drivers and co-pilots stood lined up in the same order that the 11 cars would be deployed, received the briefing, then walked out to the 11 parked cars. Next, we verified that all Raspberry Pis were running. Professor Work radioed to each car in order to ensure all vehicles and drivers were ready to go.

To estimate the time that it would take for a planned 15 minute break for the drivers at a given point of the experiment, we used the ideal spacing of 20 seconds between AVs that would load the last car onto I-24 as the first car has finished a lap. We looked into when such a break made sense to start, given the location of the Cane Ridge Event Center, a strong candidate for our Field Headquarters (FHQ). We found that it should take approximately three minutes to drive from Exit 62 to the Event Center on Old Hickory Blvd. A proposed 15 minute break was investigated.

The experimental design for the medium-scale experiment required the AVs and HDVs to move through trafic in a predefined platoon, recall Fig. 1. In order to successfully move from the staging area to the testbed as a platoon, a police escort was required to allow the entire platoon to pass through several intersections. The requirement for a police escort was one factor among many for the changes in the large-scale experimental design.

The data gathered by the eleven cars was stored on a Raspberry Pi, then uploaded via MiFi device to our CyVerse workspace. This was nontrivial to accomplish for eleven cars since they needed to be parked in a place with uninterrupted connection to upload the data. Instead, the large experiment will be captured by the I-24 Motion camera system of 294 cameras that is being installed currently, Barbour et al. (2020).

4. DRIVER INSTRUCTION AND TRAINING

Medium-scale. For the medium-scale experiment, eleven cars participated. Four of these were equipped with controllers to replace the stock ACC, and seven were manually driven. The drivers of the four AVs were part of the hardware or controller-design teams that installed and tested how the vehicles behaved. They performed these initial tests on a track and practiced using the controllers in live trafic prior to the experiment.

Large-scale. The week before the experiment, the drivers will be organized into six cohorts (of 16 or 17 drivers each). To train drivers, we will train one cohort at a time. Cohorts will be shown their driving route, and take their car to drive either the Green route or the Purple route. They will practice the procedure to exit the storage parking lot, and they will drive one loop of their respective route so that they are familiar on test day where the need to drive. They will visit the Field Headquarters site so that they know how to take their break for a driver switch, then finally practice their be guided through the parking procedure. All of these tasks are done with the same entrance and exits to the lot to simplify instructions to drivers.

ACKNOWLEDGEMENTS

The authors would like to acknowledge collaborators Eugene Vinitsky and Nathan Lichtlé for discussion and clarification of their controller algorithm development, as well as Fangyu Wu, Eugene Vinitsky, Zhe Fu, Yanbing Wang, Nour Khoudari, Joy Carpio, Fadwa Alaskar, and Kathy Jang for their assistance in data post-processing.

This material is based upon work supported by the National Science Foundation under Grant Numbers CNS-1837244 (A. Bayen), CNS-1837652 (D. Work), CNS-1837481 (B. Piccoli), CNS-1837210 (G. Pappas), CNS-1446715 (B. Piccoli), CNS-1446690 (B. Seibold), CNS-1446435 (J. Sprinkle, R. Lysecky), and CNS-1446702 (D. Work); also Grant Numbers DBI-0735191, DBI-1265383, and DBI-1743442 (CyVerse; URL: www.cyverse.org). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

This material is trial upon work supported by the U.S. Department of Energy's Ofice of Energy Eficiency and Renewable Energy (EERE) award number CID DE-EE0008872. The views expressed herein do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

REFERENCES

Barbour, W., Gloudemans, D., Cebelak, M., Freeze, P.B., and Work, D.B. (2020). Interstate 24 motion open road testbed. URL https://i24motion.org.

Bayen, A.M., Lee, J.W., Piccoli, B., Seibold, B., Sprinkle, J.M., and Work, D.B. (2020). CIRCLES: Congestion Impacts Reduction via CAV-in-the-loop Lagrangian Energy Smoothing. Presented at the 2020 Vehicle Technologies Ofice Annual Merit Review, Washington, DC (virtual).

Bayen, A.M., Lee, J.W., Piccoli, B., Seibold, B., Sprinkle, J.M., and Work, D.B. (2021). CIRCLES: Congestion Impacts Reduction via CAV-in-the-loop Lagrangian Energy Smoothing. Presented at the 2021 Vehicle Technologies Ofice Annual Merit Review, Washington, DC (virtual).

Bhadani, R., Bunting, M., Nice, M., Tran, N.M., Elmadani, S., Work, D., and Sprinkle, J. (2022). Strym: A python package for real-time CAN data logging, analysis and visualization to work with USB-CAN interface. In *2022 2nd Workshop on Data-Driven and Intelligent Cyber-Physical Systems for Smart Cities Workshop (DI-CPS)*, 14–23. IEEE.

Bhadani, R., Piccoli, B., Seibold, B., Sprinkle, J., and Work, D.B. (2018). Dissipation of emergent trafic waves in stop-and-go trafic using a supervisory controller. In *2018 IEEE Conference on Decision and Control (CDC)*, volume 57. IEEE, Fontainbleau, Miami Beach, USA. doi:10.1109/CDC.2018.8619700. URL https://ieeexplore.ieee.org/document/8619700.

Bunting, M., Bhadani, R., and Sprinkle, J. (2021). Libpanda: A high performance library for vehicle data collection. In *Proceedings of the Workshop on Data-Driven and Intelligent Cyber-Physical Systems*, DI-CPS'21, 32–40. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3459609.3460529. URL https://doi.org/10.1145/3459609.3460529.

Chou, F.C., Bagabaldo, A.R., and Bayen, A.M. (2022). The lord of the ring road: a review and evaluation of autonomous control policies for trafic in a ring road. *ACM Transactions on Cyber-Physical Systems (TCPS)*, 6(1), 1–25.

Elmadani, S., Nice, M., Bunting, M., and Bhadani, R. (2021). From CAN monitoring and data recording bridge. In *Proceedings of the Workshop on Data-Driven and Intelligent Cyber-Physical Systems*, DI-CPS'21, 17–21. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3459609.3460531. URL https://doi.org/10.1145/3459609.3460531.

Gloudemans, D., Barbour, W., Gloudemans, N., Neuendorf, M., Freeze, B., ElSaid, S., and Work, D.B. (2020). Interstate-24 motion: Closing the loop on smart mobility. In *2020 IEEE Workshop on Design Automation for CPS and IoT (DESTION)*, 49–55. IEEE.

Herrera, J.C., Work, D.B., Herring, R., Ban, X.J., Jacobson, Q., and Bayen, A.M. (2010). Evaluation of trafic data obtained via gps-enabled mobile phones: The mobile century field experiment. *Transportation Research Part C: Emerging Technologies*, 18(4), 568–583.

Jin, I.G., Avedisov, S.S., He, C.R., Qin, W.B., Sadeghpour, M., and Orosz, G. (2018). Experimental validation of connected automated vehicle design among human-driven vehicles. *Transportation research part C: emerging technologies*, 91, 335–352.

Lee, J.W., Gunter, G., Ramadan, R., Almatrudi, S., Arnold, P., Aquino, J., Barbour, W., Bhadani, R., Carpio, J., Chou, F.C., Gibson, M., Gong, X., Hayat, A., Khoudari, N., Kreidieh, A.R., Kumar, M., Lichtlé, N., McQuade, S., Nguyen, B., Ross, M., Truong, S., Vinitsky, E., Zhao, Y., Sprinkle, J., Piccoli, B., Bayen, A.M., Work, D.B., and Seibold, B. (2021). Integrated framework of vehicle dynamics, instabilities, energy models, and sparse flow smoothing controllers. In *Proceedings of the Workshop on Data-Driven and Intelligent Cyber-Physical Systems*, DI-CPS'21, 41–47. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3459609.3460530. URL https://doi.org/10.1145/3459609.3460530.

Lichtlé, N., Vinitsky, E., Nice, M., Seibold, B., Work, D., and Bayen, A.M. (2022). Deploying trafic smoothing cruise controllers learned from trajectory data. In *2022 International Conference on Robotics and Automation (ICRA)*, 2884–2890. IEEE.

Ma, J., Li, X., Shladover, S., Rakha, H.A., Lu, X.Y., Jagannathan, R., and Dailey, D.J. (2016). Freeway speed harmonization. *IEEE Transactions on Intelligent Vehicles*, 1(1), 78–89.

Nguyen, B. and Lee, J.W. (2020). CIRCLES Consortium. URL https://circles-consortium.github.io/.

Sprinkle, J.M., Bayen, A.M., Lee, J.W., Piccoli, B., Seibold, B., and Work, D.B. (2022). CIRCLES: Congestion Impacts Reduction via CAV-in-the-loop Lagrangian Energy Smoothing. Presented at the 2022 Vehicle Technologies Ofice Annual Merit Review, Washington, DC.

Stern, R.E., Cui, S., Delle Monache, M.L., Bhadani, R., Bunting, M., Churchill, M., Hamilton, N., Pohlmann, H., Wu, F., Piccoli, B., et al. (2018). Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments. *Transportation Research Part C: Emerging Technologies*, 89, 205–221.

Work, D.B., Tossavainen, O.P., Jacobson, Q., and Bayen, A.M. (2009). Lagrangian sensing: trafic estimation with mobile devices. In *2009 American Control Conference*,