Deploying a Model Predictive Traffic Signal Control Algorithm - A Field Deployment Experiment Case Study

Preprint

Qichao Wang,[1] Joseph Severino,[1] Harry Sorensen,[1] Jibonananda Sanyal,[1] Juliette Ugirumurera,[1] Chieh (Ross) Wang,[2] Andy Berres,[2] Wesley Jones,[1] Airton Kohls,[3] and Rajesh Paleti Ravi VenkataDurga[2]

1 National Renewable Energy Laboratory
2 Oak Ridge National Laboratory
3 University of Tennessee

Presented at 25th IEEE International Conference on Intelligent Transportation Systems (IEEE ITSC 2022)
Macau, China
October 8-12, 2022

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-2C00-83322
October 2022

Deploying a Model Predictive Traffic Signal Control Algorithm - A Field Deployment Experiment Case Study

Preprint

Qichao Wang,[1] Joseph Severino,[1] Harry Sorensen,[1] Jibonananda Sanyal,[1] Juliette Ugirumurera,[1] Chieh (Ross) Wang,[2] Andy Berres,[2] Wesley Jones,[1] Airton Kohls,[3] and Rajesh Paleti Ravi VenkataDurga[2]

1 National Renewable Energy Laboratory
2 Oak Ridge National Laboratory
3 University of Tennessee

Suggested Citation
Wang, Qichao, Joseph Severino, Harry Sorensen, Jibonananda Sanyal, Juliette Ugirumurera, Chieh (Ross) Wang, Andy Berres, Wesley Jones, Airton Kohls, and Rajesh Paleti Ravi VenkataDurga. 2022. *Deploying a Model Predictive Traffic Signal Control Algorithm - A Field Deployment Experiment Case Study: Preprint*. Golden, CO: National Renewable Energy Laboratory. NREL/CP-2C00-83322. https://www.nrel.gov/docs/fy23osti/83322.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-2C00-83322
October 2022

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Vehicle Technologies Office. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Deploying a Model Predictive Traffic Signal Control Algorithm - A Field Deployment Experiment Case Study

Qichao Wang[1], Joseph Severino[1], Harry Sorensen[1], Jibonananda Sanyal[1], Juliette Ugirumurera[1], Chieh (Ross) Wang[2], Andy Berres[2], Wesley Jones[1], Airton Kohls[3], and Rajesh Paleti Ravi VenkataDurga[2]

Abstract— **This paper presents a field deployment experiment of a real-time traffic signal control algorithm. We implemented the model predictive control (MPC) algorithm based on the virtual phase-link (VPL) model. We selected the deployment locations and times based on an energy saving potential concept. We developed a set of experiment systems, which included sensing, processing, and actuating components, to enable field deployment. We tested the systems rigorously before the experiment days. We reported the key procedures on the experiment days, including the steps taken, the real-time control procedure, and the monitoring of the experiment. We evaluated the impact of the deployment by looking at the changes in delay and energy consumption.**

I. INTRODUCTION

There is a rich body of literature on real-time traffic signal control methodologies. Most of these methodologies are theoretically derived and have been tested in simulations. Other than SCOOT [1], SCATS [2], and ACS-lite [3], not many real-time traffic signal control algorithms have been widely deployed in the field. Although research and development of novel real-time traffic signal control algorithms largely originates from academia or research institutes, the deployment of these algorithms is usually performed by industry. Deployment of research outcomes is critical to bridging the gap between traffic science and traffic engineering.

Compared with a simulation environment, the real world presents more challenges with implementing a real-time traffic signal control algorithm. Unlike simulations, where the ground truth is known, the real world may miss certain critical sensor data or have low-quality sensor data. In addition, the algorithm may not be executed reliably, and the communication between the systems that execute the algorithm and the field controllers may not be reliable. Beyond the technical challenges, coordination with local jurisdictions and a fail-safe considering different levels of worst condition are also critical.

The purpose of this paper is not to argue that academia and research institutes should take over industry work in deployment, but to encourage more researchers to test their algorithms in the field by easing deployment difficulties. This would allow academic research to address more field issues based on deployment outcomes and would further encourage more deployable research, yielding a greater and broader impact. We aim to provide an example of a real-time traffic signal control field deployment experiment and share the experience that we gained.

In this work, we chose to deploy the virtual phase-link-based (VPL-based) model predictive control (MPC) algorithm [4] because it fits the control area nicely. In a previous work, we modified the VPL-based traffic signal control algorithm for offline signal timing and deployed the optimized timing in Chattanooga, Tennessee [5]. In this work, we implemented the original online version of the VPL-based MPC algorithm.

In this rest of this paper, we will cover the activities before, during, and after the deployment experiment days (in the Pre-Experiment Preparations section, Experiment Days section, and Results and Analysis section, respectively).

II. PRE-EXPERIMENT PREPARATIONS

The majority of the work for the deployment experiment was done in preparation for the experiment. Before the experiment days, we selected the experiment location and time spans, understood the existing field setup, developed the experiment systems, and tested different elements of the systems rigorously.

A. Experiment Location and Time Selection

The goal of the deployed control algorithm was to improve mobility and energy efficiency. We gathered local knowledge from Chattanooga Department of Transportation (CDOT) engineers to identify the locations that needed mobility improvements. We also developed the concept of traffic-related energy saving potential.

We used energy loss, i.e., the difference between the expected vehicle-generated energy consumption per vehicle and the vehicle-generated energy consumption under free flow condition per vehicle, as a surrogate of the traffic-related energy saving potential. The vehicle-generated energy consumption is impacted by the vehicle speed, acceleration, and grade. We used a macroscopic fuel consumption model developed at NREL, RouteE [6], to estimate the vehicle-generated energy consumption. We used the information generated from probe data to universally estimate the vehicle-generated energy consumption over the whole city of Chattanooga. The probe-data-generated traffic state information in this study was acquired from TomTom.

*This work was supported by the U.S. Department of Energy (DOE).

[1]Qichao Wang (Qichao.Wang@nrel.gov), Joseph Severino, Harry Sorensen, Juliette Ugirumurera, Jibonananda Sanyal, and Wesley Jones are with Computational Science Center at the National Renewable Energy Laboratory in Golden, Colorado, USA.

[2]Chieh (Ross) Wang, Andy Berres, and Rajesh Paleti Ravi VenkataDurga are with Oak Ridge National Laboratory in Oak Ridge, Tennessee, USA.

[3]Airton Kohls is with the University of Tennessee – Knoxville in Knoxville, Tennessee, USA.

Figure 1 shows the heat map of the energy saving potential of the Chattanooga region. Colors closer to red indicate greater energy potential, and colors closer to blue indicate lower energy potential. We can see that, in addition to the downtown Chattanooga area, the area around the Shallowford Road arterial also had high energy saving potential. The Shallowford Road area was also identified by the CDOT engineers as one of the most congested areas in Chattanooga.

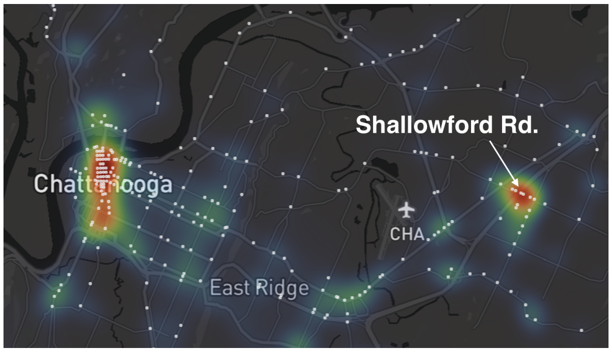

Fig. 1. A heat map showing traffic-related energy potential (colors closer to red indicate greater energy potential, and colors closer to blue indicate lower energy potential) overlaid with white dots showing where the signalized intersections are.

Figures 2 and 3 show the temporal energy loss for the eastbound and westbound traffic at the four intersections along Shallowford Road, i.e., Amin Drive and Shallowford Road, southbound ramp and Shallowford Road, northbound ramp and Shallowford Road, and Napier Road and Shallowford Road. The figures were generated with data from January 2019 to June 2021. The horizontal axis is the hour of the day, and the vertical axis is the energy loss. The grey lines represent different days, and the red lines represent the median value of the two-and-a-half years of data. We can see that both noon and afternoon have higher energy losses than other times of day.

We conducted an offline traffic signal timing experiment in February 2020 around the afternoon peak hours [5]. We chose a different time of day this time for the experiment. The noon timing in these four intersections was from 11:30 a.m. to 1:30 p.m. local time (EDT). We therefore chose the experiment to be conducted from 11:30 a.m. to 1:30 p.m. during the experiment days too address the noon hours.

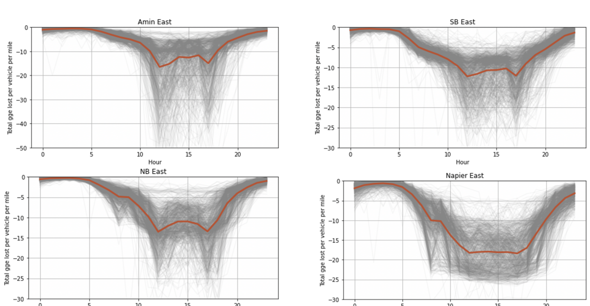

Fig. 2. Eastbound energy loss over time.

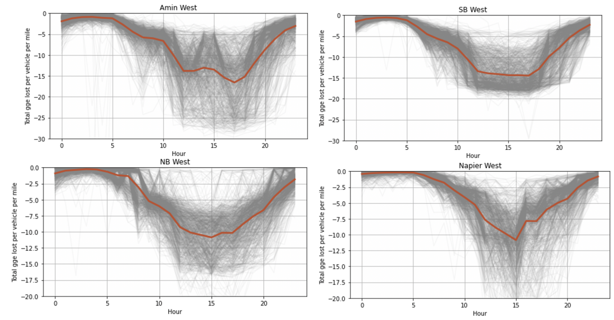

Fig. 3. Westbound energy loss over time.

B. Existing Field Setup

Figure 4 shows a map of the studied area. The four intersections that were controlled are circled on the map. From west to east, these four intersections are: Amin Drive and Shallowford Road, southbound ramp and Shallowford Road, northbound ramp and Shallowford Road, and Napier Road and Shallowford Road. We will reference the intersections by their side street names in the rest of this paper. There is a mall area at the east side of the studied area that attracts high volume of traffic. A significant amount of traffic from the mall area goes to the highway through westbound and northbound of Napier intersection. To make the traffic smoother along the arterial, CDOT engineers wanted to avoid having the queues spill back to the highway, and they also wanted the queues from the mall to be managed at a reasonable level.

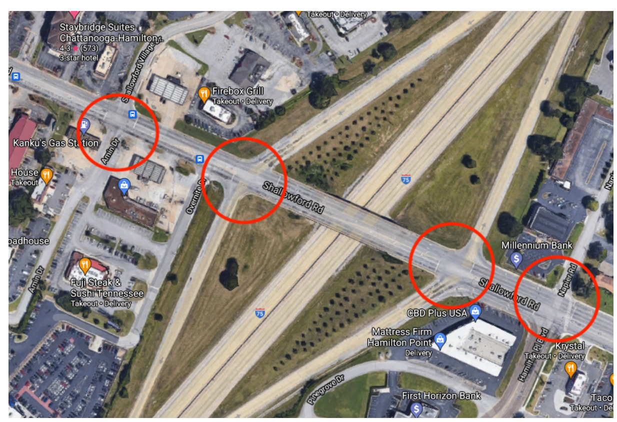

Fig. 4. Shallowford Road map.

These intersections were controlled by Siemens M60 controllers in TS2 Type 2 cabinets. The M60 controllers were in version 3.x. To support the real-time signal control experiment, CDOT upgraded the controllers to version 5.x. CDOT engineers centrally configured the signal controllers in the region using Siemens' TACTICS software. Traffic signal controllers in the United States use the National Transportation Communications for ITS (Intelligent Transportation Systems) Protocol (NTCIP) as an industry communication protocol. To ensure reliable and scalable real-time traffic signal control, the Siemens Mobility team provided the research team with their NTCIP server. The NTCIP server receives and sends out commands to the targeting traffic signal controllers in a reliable and scalable manner.

The four intersections were equipped with GridSmart cameras. A GridSmart camera is a fish-eye camera installed above an intersection that can actuate traffic signals and provide sensed traffic state information. The GridSmart camera system can report the traffic volume in each movement direction.

The studied area also has probe vehicle information coverage through TomTom. TomTom provided historical hourly aggregated link-level speed, travel time, and volume estimation information. It can also provide this information in real time through its REpresentational State Transfer (REST) application programming interface (API), but at a lower spatial resolution.

C. Experiment Systems

We developed software systems to support the field deployment. Figure 5 shows the architecture of the experiment systems, which have sensing, processing, and actuating parts.

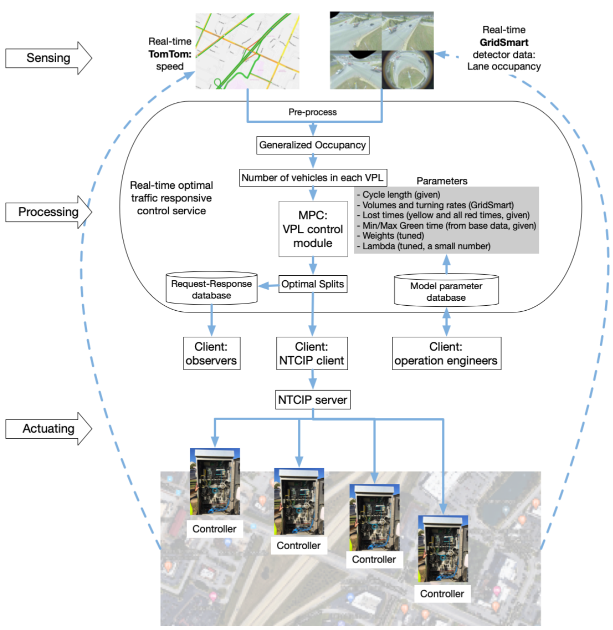

Fig. 5. Real-time traffic signal control field experiment system architecture.

1) Sensing: We used real-time traffic state data from both stationary sensors (GridSmart cameras) and probe data (TomTom). We used the stationary sensor data mainly for the better resolution and its ability to provide lane-based information. The lane-based information is crucial to distinguishing left turn movement from through movement going westbound on the bridge. It was observed that, on the bridge, the westbound left turn movement did not have enough green time and caused queue spill-back, whereas the westbound through movement had more than enough green time.

Fish-eye cameras only have good field of view near the area under the camera and do not have good sensing in the area farther away. However, the level of congestion is represented better by the upper stream of the road link than the stop bar area. We therefore placed virtual advanced detectors in the upstream intersection camera and looked toward the downstream intersection to obtain the real-time congestion level of the incoming link of the down stream intersection (Figure 6).

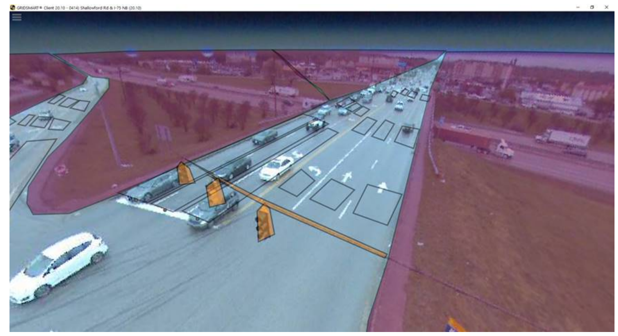

Fig. 6. Virtual detectors in a GridSmart camera view.

We placed multiple virtual detectors on each lane so that we could choose the best-fit detector data to inform control. We picked the detectors whose data could form well-structured fundamental diagrams. We used occupancy from the advanced detector to reflect the density of the link. Figures 7 and 8 show the volume vs. occupancy plots, aggregated every 5 minutes, for two advanced detectors from the GridSmart camera at the northbound ramp intersection. We selected the detector from Figure 8 over the detector from Figure 7 because it formed a better structured fundamental diagram and could clearly show different levels of congestion.

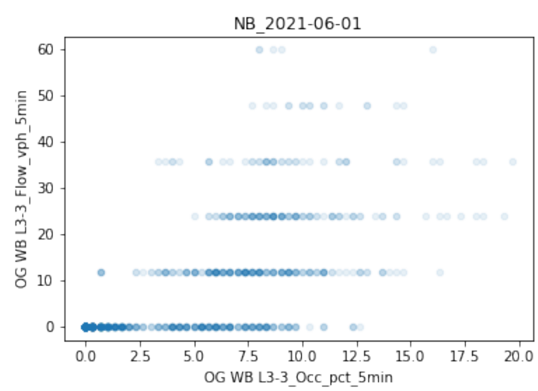

Fig. 7. Volume vs. occupancy from a bad detector.

For the virtual phase links without good GridSmart detector data, we assumed constant numbers to reflect the level of congestion based on historical observations. We were able to use the ramp speeds from the real-time TomTom REST API to estimate the real-time congestion levels at the northbound and southbound ramps. We could not derive other road links' congestion levels from the real-time TomTom REST API because the other links from the TomTom real-time information map covered multiple intersections.

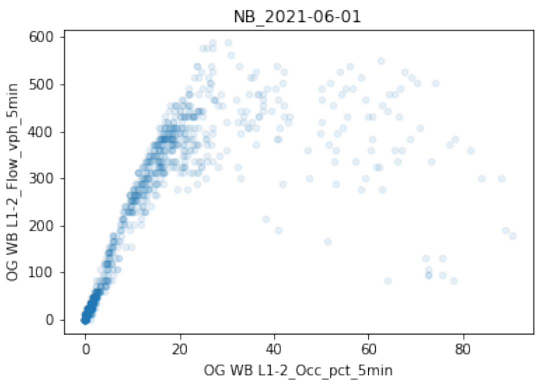

Fig. 8. Volume vs. occupancy from a good detector.

2) Processing: The algorithms for processing was wrapped in a docker container and hosted on the cloud. The new, optimized timing can be requested over the public internet with an authorized API key. Both real-time TomTom data and real-time GridSmart data were preprocessed and converted into generalized occupancy information and then converted into the number of vehicles in each VPL. The real-time GridSmart API provided minute-by-minute information for the past hour. We averaged the past five minutes information to avoid oscillation in sensing and control. The real-time TomTom API provided only the information about the current minute. We therefore stored the real-time TomTom data in a database every minute through a cron (a time-based job scheduler) job. When the VPL-based MPC module requests TomTom data, it receives the data averaged over the previous five minutes. The core algorithm for the processing part is the model predictive control (MPC) algorithm. The input to the MPC algorithm is the number of vehicles in each VPL. The MPC algorithm has a set of parameters that need to be set to produce an optimal timing. Most parameters were either given by CDOT engineers or could be reported by GridSmart or TACTICS. The other parameters could be tuned as we deployed the systems. All the model parameters were saved in a database, which could be updated via an API with authorized API keys. The optimal splits were returned to the authorized requester as responses. All the requests and responses were saved to a database that can be visualized via a dashboard that we developed to monitor the experiment.

3) Actuating: The new timing was ingested into the field controllers through NTCIP. Siemens Mobility provided their NTCIP server to support more reliable and scalable communication to the field controllers. We developed a client for the NTCIP server that would automatically pull the latest optimal timing, verify if the timing was legit to the controllers, and push the latest timing to the NTCIP server. After the end of the experiment, we used the client to set back the existing timing before the experiment to the field controllers.

D. Testing

Before the field deployment, we rigorously tested the systems. We tested the real-time data stream, real-time optimal traffic signal timing service on the cloud, and the NTCIP server connection.

The GridSmart data can only be accessed within the CDOT network. The NTCIP server was also hosted within CDOT. We tested the VPN connection and the functionalities of the NTCIP server.

For the real-time control MPC algorithm, we developed local-sandbox-staging three level tests. We started with testing the system at staging on the cloud. If we passed the staging test, we deployed the algorithms in production. If an error occurred, we pushed the code to a sandbox to identify the error message. (This is because the staging service does not show the error message.) If more debugging was needed, we moved on the test the code locally. We tested each module of the algorithm code to identify any bugs.

III. Experiment Days

The field deployment experiment was conducted for noon peak hours (11:30 a.m. to 1:30 p.m.) from June 8, 2021, to June 11, 2021. We contained the real-time control to be within this time of day and only changed the splits in the timing plan for the noon peak hours. The research team met virtually with CDOT engineers 15 minutes before the start of the experiment (i.e., 11:15 a.m.). During this time, we changed the signal controllers of interest to ACS-lite mode through TACTICS software and initialized the experiment system. The existing timings were recorded before the start of the real-time control and were changed back after the experiment time.

Figure 9 illustrates the cyclic real-time control procedure. The procedure starts with measuring the current traffic states. The traffic state signal is smoothed with a moving average of the past five minutes. The measured current states are then fed into the algorithm to optimize the traffic signal timings. (These optimized timings were implemented through the NTCIP server.) The systems then wait for k minutes ($k = 4$ in this experiment) before starting the next cycle of measuring states, optimizing timings, and implementing the optimized timings.

The experiments were monitored by the research team and partners from CDOT. Figure 10 shows a screenshot of the terminal information and camera views. The top four terminals were directly connected to the field signal controllers through secure shell (SSH) protocol. These four terminals were operated the same as the control panels on the field signal controllers. Under the four signal controller terminals, we placed the camera views from the fish-eye cameras above each controlled intersection (i.e., GridSmart cameras). The bottom left terminal ran the code to periodically pull the latest optimal timing from the cloud and send the formatted information to the NTCIP server. The terminal on the bottom right showed the NTCIP server outputs (e.g., the information received from the code running in the terminal on the left side). We also placed the view of a SmartWay camera looking at the southbound off-ramp to make sure the queue did not spill back onto the highway.

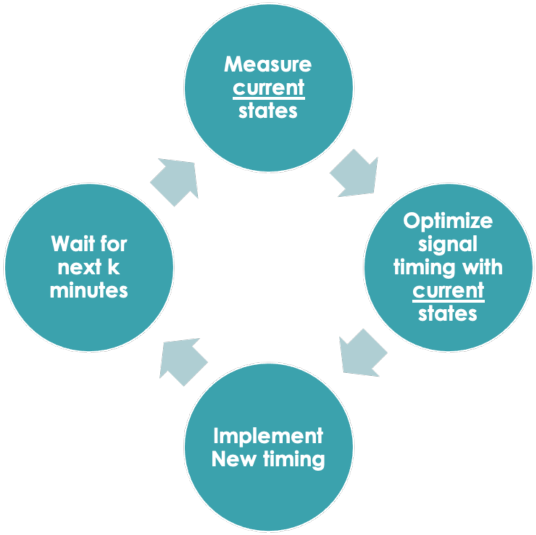

Fig. 9. Real-time control procedure.

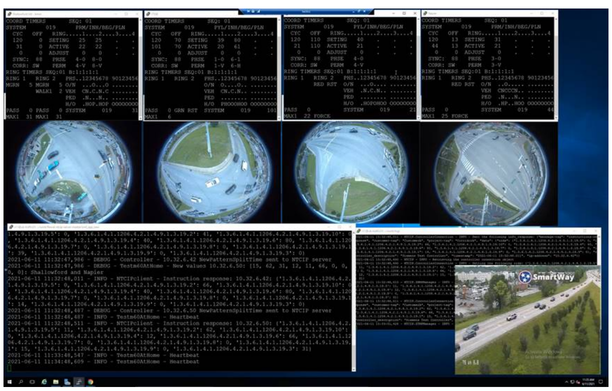

Fig. 10. Screenshot of experiment monitoring through cameras and terminals.

Besides monitoring what was happening in the field, we also developed a dashboard to show what the sensors think are happening in the field. The real-time sensor information was monitored because the control algorithm was fed by the real-time sensor data. Figure 11 shows part of the dashboard that plotted the real-time TomTom data.

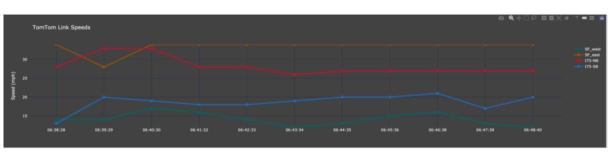

Fig. 11. Experiment monitoring dashboard.

IV. Results and Analysis

We evaluated the impact of the control deployment by comparing the travel delays and energy consumption during the experiment with the performance before and after the experiment.

Usually, evaluating the performance of a field deployment involves looking at a longer period of time. We evaluated the traffic demand in the study area. Figure 12 shows the estimated volume distribution in the past two years from the experiment time. The volume was estimated using TomTom data with a machine learning model from [7].

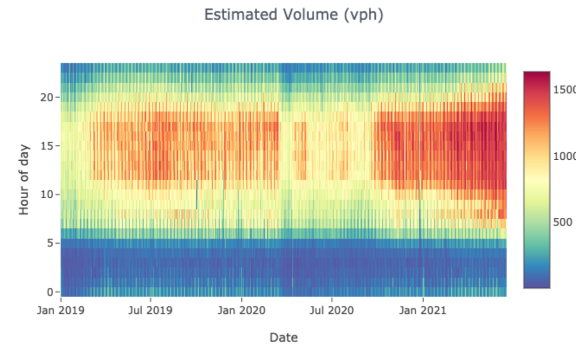

Fig. 12. Historical volume at Shallowford Road.

In Figure 12, the horizontal axis shows the date, from January 2019 to June 2021. The vertical axis shows the hour of day, from midnight to midnight. The colors represent the estimated volumes from the historical TomTom data. As we can observe from the figure, the volume dropped to a low level during the peak of COVID-19 shutdowns in 2020. The volume increased above the pre-COVID average since the end of 2020. The traffic demand pattern was still changing around the time of the experiment. Therefore, we only compared the data from the experiment with the week before and after the experiment.

We evaluated two metrics: change in delay and change in energy consumption. We calculated the delay by calculating the difference between the travel time and the free flow travel time. The free flow travel time was approximated by the tenth percentile travel time from the historic data. The average vehicle energy consumption $e_{i,t}$ was computed from the energy pipeline that was developed at NREL using TomTom data. The detailed description of the energy pipeline was presented in [8].

Table I shows the percentage change in delay compared to the week before the experiment over the four intersections on Shallowford Road. In the evaluation, we include both the main road along Shallowford Road and the minor roads perpendicular to the corridor. The experiment was run for two hours beginning on a Tuesday at 11:30 a.m. EST. The data we used to evaluate the experiment was originated from TomTom. The TomTom historic data was provided to us as hourly aggregates. Thus, hour 11 (i.e., 11:00 a.m.) was controlled for half of the hour, and the other half-hour was under default settings. We can see from Table I that the delay was improved during hour 12, showing a significant improvement to the flow of the corridor. The real-time control was best on Tuesday and Wednesday, improving the delay by more than 30%. The least improvement was on Friday, June 11, when the delay improved by just 8%. Looking at the 11th hour, we still see significant improvements to delay, but unfortunately, we could not see the full actualization of the real-time control improvement because half of the hour was controlled under default settings.

TABLE I
PERCENTAGE CHANGES IN TOTAL DELAY FROM THE WEEK BEFORE THE EXPERIMENT

Hour	June 8	June 9	June 10	June 11
11	-17.67%	-9.48%	-16.25%	4.15%
12	-30.17%	-30.07%	-22.59%	-7.89%

Looking now at Table II, we continue to see reduced delays along the corridor, with Wednesday, June 9, at hour 12 showing the greatest improvement. The least improvement in delay was on Friday, much like our comparison to the week before. However, the delay improved by more than 17% compared to the following Friday. We again see that hour 11 had less of an improvement than hour 12, but this was expected, because we only controlled for half of the time in that hour. The 11th hour delay metrics were all improvements over our comparison week.

TABLE II
PERCENTAGE CHANGES IN TOTAL DELAY FROM THE WEEK AFTER THE EXPERIMENT

Hour	June 8	June 9	June 10	June 11
11	-7.36%	-18.82%	-4.38%	-9.17%
12	-20.01%	-32.16%	-23.13%	-17.13%

Doing the same analysis as above, but using total energy as our comparison metric, we found similar improvements to the week before and the week after. Tables III and IV show improvements ranging from 1% to nearly 5% during the 12th hour of experiment. Again, we see that, on Wednesday and Friday, the 11th hour metrics are less improved—or even worse—than the week before (see Table III). Additionally, the delay metrics show much greater improvements than the percentages in energy metrics. This is due to the algorithm's objective to maximize flow and minimize delay without accounting for energy. In the future, we will be using reinforcement learning with rewards based on minimizing total energy consumption.

We also acquired high-resolution probe vehicle trace data (i.e., Wejo data) and derived the travel time from the probe vehicle trace data. For those four days, we reduced overall delay by 45%. The delay was improved by 57%, 61%, 47% and 14% respectively through the week.

TABLE III
PERCENTAGE CHANGES IN TOTAL ENERGY FROM THE WEEK BEFORE THE EXPERIMENT

Hour	June 8	June 9	June 10	June 11
11	-0.77%	1.05%	-3.48%	0.01%
12	-4.76%	-2.89%	-3.86%	-1.11%

TABLE IV
PERCENTAGE CHANGES IN TOTAL ENERGY FROM THE WEEK AFTER THE EXPERIMENT

Hour	June 8	June 9	June 10	June 11
11	-1.58%	-2.05%	-1.94%	-2.29%
12	-3.03%	-3.98%	-4.62%	-4.14%

V. FURTHER WORK

One obstacle that makes the transfer of in-house traffic science research to real-world deployment difficult is that the real-world applications have less observability than the simulation. We therefore built a digital twin to link simulations with real-world deployment.

We used SUMO traffic simulation software to model and simulate the study area. We obtained the traffic volume and turning ratio from the GridSmart camera data and obtained the signal timing from the Chattanooga Department of Transportation. The existing SUMO package does not support dual-ring National Electrical Manufacturers Association (NEMA) timing, so we developed a signal control module in SUMO that supports NEMA timing [9]. Beyond traffic demand and signal control, we also modeled the real-time traffic state sensors in the SUMO simulation. We used lane area detectors in SUMO to simulate the virtual advanced detectors in GridSmart cameras. We matched the detectors' names and locations to the ones in the GridSmart cameras' settings. We also modeled the TomTom data feed by getting probe vehicle speeds from the simulated network links that corresponded to the TomTom links.

VI. CONCLUSIONS

In this paper, we reported an experiment that deployed a real-time traffic signal control algorithm. It demonstrated the elements and steps involved in deploying such an algorithm as well as the challenges with real-world deployment. The intention of this paper is to share the experience of a field deployment experiment to ease the difficulties for researchers who are deploying their algorithms.

This paper can be used as an example of deploying a real-time traffic signal control algorithm. Although we implemented a specific algorithm, i.e., VPL-based MPC, the steps and system settings can also be utilized for other real-time traffic signal control algorithms.

The reported experiment was conducted during the noon peak hours. More work needs to be done to allow the algorithm run 24/7.

VII. ACKNOWLEDGEMENTS

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Vehicle Technologies Office. A portion of The research was performed using computational resources sponsored by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

This work was made possible through the close cooperation of the Chattanooga Department of Transportation and Siemens Mobility. The research team acknowledges and appreciates guidance and technical support of Cindy Shell, Tommy Trotter, Casey Lewis, and Michael Gaertner.

References

[1] P. Hunt, D. Robertson, R. Bretherton, and R. Winton, "SCOOT-a traffic responsive method of coordinating signals," tech. rep., 1981.

[2] P. Lowrie, "SCATS, sydney co-ordinated adaptive traffic system: A traffic responsive method of controlling urban traffic," 1990.

[3] F. Luyanda, D. Gettman, L. Head, S. Shelby, D. Bullock, and P. Mirchandani, "ACS-lite algorithmic architecture: applying adaptive control system technology to closed-loop traffic signal control systems," *Transportation Research Record*, vol. 1856, no. 1, pp. 175–184, 2003.

[4] Q. Wang and M. Abbas, "Optimal urban traffic model predictive control for NEMA standards," *Transportation Research Record*, vol. 2673, no. 7, pp. 413–424, 2019.

[5] Q. Wang, J. Severino, J. Ugirumurera, W. Jones, and J. Sanyal, "Offline arterial signal timing optimization for closely spaced intersections," in *2021 IEEE Green Technologies Conference (GreenTech)*, pp. 344–350, IEEE, 2021.

[6] J. Holden, N. Reinicke, and J. Cappellucci, "RouteE: A vehicle energy consumption prediction engine," *Society of Automotive Engineers Technical Paper Series*, vol. 2, no. NREL/JA-5400-78089, 2020.

[7] Y. Hou, V. Garikapati, C. Hoehne, K. Kasundra, and S. Young, "Statewide traffic volume estimation for non-freeway roads using probe-vehicle data and machine learning methods," tech. rep., 2021.

[8] J. Severino, Y. Hou, A. Nag, J. Holden, L. Zhu, J. Ugirurmurera, S. Young, W. Jones, and J. Sanyal, "Real-time highly resolved spatial-temporal vehicle energy estimation using machine learning and probe data," tech. rep., 2021.

[9] Q. Wang, W. Jones, and T. Li, "NEMA-Phase-Compliant Traffic Signal Controller Module in SUMO," 2021.