# A Vision of Smart Traffic Infrastructure for Traditional, Connected, and Autonomous Vehicles

Sanjay Ranka, Anand Rangarajan, Lily Elefteriadou, and Siva Srinivasan
University of Florida, Gainesville, FL

Emmanuel Poasadas and Dan Hoffman
City of Gainesville
Gainesville, FL

Raj Ponnulari, Jeremy Dilmore, and Tom Byron
Florida Dept. of Transportation
Tallahasee, Florida

**Abstract — This smart city traffic management approach seeks to use edge-based video-stream processing (using multicore and GPU processors) at intersections and in public vehicles (city buses, fire trucks, ambulances, school buses) to convert video data into space-time trajectories of individual vehicles and pedestrians that are transmitted to a cloud-based system. Key information is then synthesized in the cloud from them to create a real-time city-wide traffic palette. Real-time or offline processing both at the edge and the cloud will then be leveraged to optimize intersection operations, manage network traffic, identify near-collisions between various units of traffic, provide street parking information, and a host of other applications. Additional information such as weather and environment will also be leveraged.**

**The use of edge-based real-time machine learning (ML) techniques and videostream processing has several significant advantages. (1) Because there is no need to store copious amounts of video (few minutes typically suffice for edge-based processing), it automatically addresses concerns of public agencies who do not want person-identifiable information to be stored for reasons of citizen privacy and legality. (2) The processing of the video stream at the edge will allow for the use of low bandwidth communication using wireline and wireless networks to a central system such as a cloud, resulting in a compressed and holistic picture of the entire city. (3) The real-time nature of processing enables a wide variety of novel transportation applications at the intersection, street, and system levels that were not possible hitherto, significantly impacting safety and mobility.**

***Keywords—video processing, sensor data, machine learning, big data analytics, intelligent transportation, smart city***

## 1. Introduction

Mitigating traffic congestion and improving safety are the important cornerstones of transportation for smart cities. Despite significant advances in vehicle technology, traffic engineering practices, and analytics based on crash data, the number of traffic crashes and fatalities are still too many. An INRIX study[1] found that in 2017, traffic congestion resulted in nearly $305 billion in congestion costs and caused Americans to lose 97 hours per person in congestion. This costs the U.S. $87 billion annually in lost time (with the American Transportation Research Institute estimating the freight sector loss due to congestion at $75 billion annually). Many drivers are frustrated due to long (but potentially preventable) delays at intersections. Traffic signal control timing does not change in real-time based on crashes, incidents, or changes in traffic patterns and behavior. Addressing these challenges requires a thorough understanding of traffic patterns not only at intersections but on streets *and* in the overall network. Unfortunately, existing monitoring systems and decision making for this purpose have several limitations:

- Current sensors have limited capability: Vehicle loop detectors that have traditionally been deployed at intersections to detect the passage of vehicles are error-prone; have high deployment and maintenance costs; can only measure the absence or presence of vehicles passing above them; and are not always useful for observing the movements of pedestrians and scooters. When the sensors are not accurate or timely an adaptive strategy will not be effective. Video detection has great potential to improve accuracy and timeliness in the detection of vehicles, pedestrians, bicyclists, etc.
- Current software systems for traffic monitoring are fragmented and not suitable for real-time decision making: Transportation professionals are presented with a plethora of fragmented data in various systems. Existing intersection control systems do not provide reports on a real-time basis (based on the vendor), and these are given at coarse levels of granularity (for example, traffic movement counts by the hour) limiting their use and ability to make real-time changes to adapt to dynamically changing conditions. Current approaches are not readily scalable because of constraints of cost, bandwidth, and lack of integration.

Our approach uses *edge-based* video-stream processing (using multicore and GPU processors) at intersections and in public vehicles (city buses, fire trucks, ambulances, school buses) to convert video data into space-time trajectories of individual vehicles and pedestrians that are transmitted to a cloud-based system. Key information is then synthesized in the cloud from them to create a real-time

[1] https://inrix.com/scorecard/

978-1-7281-6059-7/20/$31.00 ©2020 IEEE
DOI 10.1109/MetroCAD48866.2020.00008

city-wide traffic palette. Real-time or offline processing both at the edge and the cloud can then be leveraged to optimize intersection operations, manage network traffic, identify near-misses between various units of traffic, provide street parking information, and a host of other applications. Additional information such as weather and environment will also be leveraged in the future. The use of edge-based real-time machine learning (ML) techniques and video-stream processing has several significant advantages:

- Because there is no need to store copious amounts of video (few minutes typically suffice for edge-based processing), it automatically addresses concerns of public agencies and civil libertarians who do not want person-identifiable information to be stored for reasons of citizen privacy and legality.
- The processing of the video-stream at the edge allows for the use of low bandwidth communication using wireline and wireless networks to a central system such as a cloud, resulting in a compressed and holistic picture of the entire city. This can then be used for efficient processing.
- The real-time nature of processing enables a wide variety of novel transportation applications at the intersection, street, and system levels that were not possible hitherto, significantly impacting the safety and mobility of our community.
- The use of cloud computing for processing aggregated data from multiple video and other sensors will provide information unification and the necessary on-demand computing horsepower for the large-scale simulations for system level optimizations and analytics. It will also allow for leveraging offline processing of historical data to be used in conjunction with real-time information.

We broadly describe our approach for improvements of community safety and mobility at three levels: smart intersections, smart streets, and smart systems (Sections 2, 3, and 4). We also describe the interaction of these technologies with connected and autonomous vehicles (Section 5). These assume that there exist video-based systems that can observe traffic at intersections or from publicly owned vehicles. The data collected from the above sources afford real-time measurements and decisions which will then be used as part of an iterative approach to observe traffic and pedestrian mobility, allowing for the introduction of new design concepts (such as leading pedestrian intervals), new technologies (like reaction to autonomous and emergency vehicles), and impacts of environmental improvements (lighting, landscaping, and geometric design changes) at intersections.

We have made considerable progress in implementing our vision of developing a smart infrastructure. The current status of the project is described in Section 6. In the long run, this work will lead to major improvements in the current state of practice: the underlying technology, the associated data, and the resulting policies and procedures are interlocked through a cross-disciplinary approach utilizing advances in computer science (video analytics, machine learning, sensor fusion) and traffic engineering. We believe that our approach can have a tangible impact on the USDOT goal of the Vision Zero plan to minimize, and eventually eliminate, motor vehicle-related crashes and accidents. The ability to detect and understand unsafe driving and walking conditions (measured as "near misses" rather than "crashes") and react to them in real time would be critical in further moving the needle towards the goal of Vision Zero.

## 2. Smart Intersections

Using standard and fisheye cameras mounted at intersections (and integrating weight sensors placed beneath streets), we are leveraging the state of the art in computer vision and machine learning to perform vehicle tracking (localization, tracking, turn estimation) and pedestrian monitoring at intersections. Traffic surveillance of dynamic objects, particularly vehicles on the road, has been an active research topic in past decades in the fields of computer vision and intelligent transportation systems. In the interests of real-time feedback, security, and citizen anonymity, we have elected to mainly focus on real-time (and near real-time) video processing at intersections.

Our goal at this juncture is to leverage the output of video processing and tracking approaches into efficient real-time (and near real-time) representations. Subsequently, these compressed representations drive applications such as the identification of near misses, summary of pedestrian movements, and the extraction of vehicle trajectories at intersections. Since video information is clocked and integrated with weight sensors, we anticipate real-time feedback to signal controllers resulting ultimately in improvements to intersection level signal planning (offsets, phase switches, etc.).

*The uplink of compressed information to the cloud is a side benefit resulting in much lower communication costs.* We expect this work to have a significant impact on understanding traffic behavior at an intersection that will lead to the following:

- Improved Pedestrian and Bicyclist Safety by examining the conflict points of the vehicle and pedestrian trajectories and anomalous behavior, based on time of day and day of the week. The availability of temporal profiles of "exposure data" (i.e., the volumes of pedestrians and bicyclists) from video-based systems is critical for performing comparative assessments of ped/bike safety at different locations.
- More Accurate Demand Profiles which can lead to more effective signal timing and a better understanding of the relationships between existing (loop detector) and new technologies (video).
- Better Incident and Bottleneck Management using real-time information and messaging of incidents and other events to pedestrians and vehicles.

## 3. Smart Streets

We shift perspective to include onboard video information (on city buses, emergency vehicles, and police cars). Once again, we plan to leverage best-of-breed computer vision algorithms to perform tracking of vehicles, assessment of near collisions, jaywalking, tailgating — essentially the entire gamut of live traffic-centered activities. Since the city buses (for example) are moving, tracking will have to be performed taking camera motion into account; the same applies for integration and cloud communication. Egomotion issues, for example, are not a problem since accurate GPS information can be used to determine camera coordinates extrinsically. Multiple cameras on the buses allow for better integration and more accurate estimation of vehicular tracks. Bluetooth communication from the bus to roadside sensors can also be integrated into the system if available.

Such data can also provide insights into how passengers approach or leave the transit buses at stops. All these can be used to understand pedestrian safety both during day and night times. In contrast, tracking vehicles in terms of trajectories and state of brake lights and indicator lights ahead of the bus can help understand lane-changing maneuvers undertaken before implementing a turn. The space-time trajectories of vehicles along adjacent lanes can be used to infer gap-acceptance behavior, which in turn leads to lane changes. Again, differences can be examined across different weather conditions. Additionally, the video collected can be used to determine free parking spots along the bus routes that can be communicated in real-time to the cloud and then to citizens. Furthermore, the integration of motion tracking, traffic pattern identification, etc. from video (and other sensors) using city buses etc. is fundamentally novel to the traffic engineering field. While summarization information from the segments will still be sent to the cloud, the edge-based processing will directly result in changes to signal parameters with these effects first tested in simulation platforms such as SUMO or VISSIM.

Inclusion of real-time (and near real-time) video processing conducted on city buses (and other authorized vehicles) has the following expected outcomes:

- Better Pedestrian Safety by detection of pedestrian movements at intersections and mid-blocks
- Better Resource Management by understanding usage of street parking and signage
- Better Lane and Street Sign Design by tracking indicator lights, lane-changing maneuvers undertaken in front of the bus.

## 4. Smart Network

We envisage a comprehensive effort at the network level aimed at nothing less than a complete characterization and activity recognition of important traffic patterns, the clustering of vehicles and pedestrians, responses to emergency vehicles, adjustments made during school closings, etc. We will use all summarized information (traffic and pedestrian activity patterns, dominant arterials, corridors, and sub-networks) sent to the cloud to perform network level traffic signal planning using pre-timed signal optimization and dynamic reinforcement learning while undertaking comprehensive policy testing in traffic simulation software such as SUMO and VISSIM.

*Machine learning on historical data:* Once all relevant traffic data have been sent to the cloud, we can begin analysis of historical data to perform a space-time decomposition of the network. Essentially, we seek to utilize the video and other sensor summarization information to carve up the network spatially into different arterials, corridors, and sub-networks and for relevant time periods. While the simplest approach we envisage is space-time clustering, this may be insufficient to extract long arterials. A more novel machine learning (ML) approach is to cluster segments (represented as graph edges) while enforcing graph connectivity and temporal compactness (e.g., a long section of Newberry road in Gainesville from 4-6 PM weekdays -- a typical example of a high-volume corridor at peak rush hour). From an ML perspective, this is an unsupervised learning problem in a spatiotemporal graph which we seek to decompose into relevant fragments while ensuring spatial connectivity (artery or corridor) and temporal compactness (a reasonable block of time). Once this step has been accomplished, we propose to first perform pre-timed signal optimization and subsequently global optimization for real-time and unforeseen events.

The decomposition of the network into important fragments (arterials etc.) has a huge payoff in larger networks: the decomposition facilitates network-level optimization of signal timing. Overall the sensing data from multiple intersections will enable

- Better Incident Detection for alleviating traffic backups and secondary crashes
- Better Signal Retiming for corridors by time of day and day of the week to reflect the changes in network demand
- Better System-wide Network Utilization using a global view of the entire network and potential disruptions.

## 5. Connected and Autonomous Vehicles

Connected and automated vehicles (CAVs) are expected to organically enter the traffic stream over the next decade. These vehicles have the potential to reduce traffic accidents and improve the efficiency of transportation systems. We expect a symbiotic relationship in which the information exchange between the infrastructure and the vehicle should lead to mutual benefits. We are deploying and testing connected vehicle technologies as part of our infrastructure work (more details are provided in Section 6A). By collecting, transmitting, and analyzing data, we expect to

- Improve safety by disseminating relevant information to vehicles, buses, pedestrians, and other units of traffic within the Trapezium (Figure 1)
- Enhance our understanding of CAV behavior and their interaction with conventional vehicles
- Develop better signal timing optimization algorithms which use advanced detection based on connectivity
- Enhance previously developed trajectory optimization algorithms for CAVs, which can significantly improve the efficiency of the highway network
- Develop improved tools to take advantage of CAV-related technologies and understand the reasons that

may affect market penetration.

The goal of the project is to improve travel time reliability, safety, throughput, and traveler information.

## 6. Current Status

Below, we provide a high-level overview of our work in the various aspects of smart traffic infrastructure.

### A. *I-Street and Trapezium Testbed*

The University of Florida Transportation Institute (UFTI)[2], with support from FDOT and the City of Gainesville (CoG), is developing a smart testbed on the UF campus and adjoining city streets. The testbed, called I-STREET, was established to deploy and evaluate numerous advanced technologies for conventional, connected, and autonomous vehicles. This includes the use of smart devices and sensors to develop novel applications. The overall testbed includes heavy pedestrian flow, extensive bicycle facilities, scooters, and mopeds on campus and a variety of highway facilities ranging from four-lane arterials to two-lane low-speed roadways. Gainesville is home to one of the most heavily used transit systems in Florida (RTS – http://go-rts.com), which also serves the UF campus.

Additionally, the Gainesville Trapezium is deploying and testing connected vehicle technologies and applications along four roads forming a trapezium surrounding the University of Florida main campus, as shown in Figure 4. The goal of the project is to improve travel time reliability, safety, throughput, and traveler information. Approximately 27 roadside units have been installed in this project. The roadways and intersections along and within this Trapezium and bus routes serving this area will constitute the fundamental real-world testbed for this study.

Most of the signalized intersections of interest to this study have live CCTV RTSP feeds streamed at 30 frames per second HD quality (720p or 1080p), pan, tilt, and zoom capabilities, and video detection for stop bar, presence, and traffic counting. On select intersections, multimodal video detection for pedestrians, bicycles, and vehicles (using Iteris Vantage Live, Smart Cycle, and PedTrax) and fisheye video detection with motion tracking and vehicle classification capabilities are already being installed or are planned for installation. ATC controllers can provide signal-timing history at decisecond resolution (upgrades to the latter are past the planning stages). Thus, a variety of video feeds and signal timings are available that will be used for ratification of the methods. The video cameras are addressable via a local network at every intersection, and the proposed research calls for processing these video feeds before they leave this network. Eight- or twelve-port copper Ethernet (fast and gigabit) connectivity capable of VLAN segmentation is available at each intersection, with a single-mode fiber backhaul to the Traffic Management Center. These will facilitate the transfer of data from the intersections to the cloud. About 5-10 intersections will be chosen to pilot test the smart-intersection efforts. Edge processors will be added to the video cameras present at these intersections (with the CoG and UF jointly designing and testing them to ensure effective real-time processing).

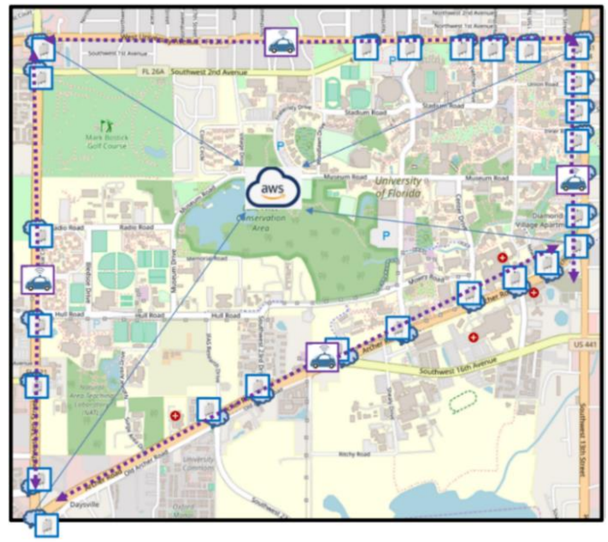


**Figure 1: Trapezium project**: This project is deploying and testing connected vehicle technologies and applications along four roads forming a trapezium surrounding the University of Florida main campus. All intersections will have roadside units to interact with CAVs. Additionally, video cameras will be installed at a large subset of these intersections.

Intelligent transportation systems require the use of interactions between road users and infrastructure. Dedicated short-range communications (DSRC) using radio, Wi-Fi, or cellular technologies can enable such interactions at signalized interactions. By using DSRC effectively, the infrastructure systems can provide information to the users about the interactions as well as using the road users as probes to create a vignette of local and network level traffic patterns and usages.

### B. *Computing and Software Architecture*

The overall approach seeks to use edge-based video-stream processing (using multicore and GPU processors) at intersections and in public vehicles to convert video data into space-time trajectories of individual vehicles and pedestrians that are transmitted and synthesized on a cloud-based system. Improving image processing techniques for traffic surveillance is an ongoing area of research in the computer vision and intelligent transportation systems (ITS) communities. For these reasons, edge computing is centered around the concept of adding a layer of compute resources between the "IoT" devices and servers on the "cloud". The goal of the intermediate layer is to improve performance (by

[2] https://www.transportation.institute.ufl.edu/research-2/istreet-about-us/

reducing network latency) and reduce the volume of data being transported (in an open and interoperable way). The value of this approach is also in enabling seamless interoperability between the heterogeneous datasets and providing the computing power necessary for bulk data processing and providing real-time applications that only need data collected locally.

*Edge Processing*: Below, we describe key computing and software challenges that we are addressing for edge video processing:

- Video Characteristics: At 20-30 frames-per-second, high-definition video from a single camera can generate over 2 gigabytes per hour. We are using GPU processing for real-time usage. Some of the main challenges faced by traffic video are the sensitivity to variations in lighting and weather and the occlusion of vehicles.
- Deep Neural Network Software: Deep neural networks (DNNs) are currently widely used (Abadi et al., 2016) for many artificial intelligence (AI) applications, by gathering knowledge from experience with a hierarchy of concepts. We have developed DNN approaches for vehicle detection and tracking.
- GPU Processing: Most image processing and deep learning frameworks can effectively leverage Nvidia's cuDNN and other libraries for rapid execution. This acceleration is transparent to the user of the framework library. Similar processing capabilities are also available on Amazon and Google cloud platforms.

*Cloud Computing:* Cloud computing platforms are available from a variety of vendors such as Google, Amazon, IBM, and Microsoft. We are using AWS from Amazon because of our experience and the fact that it provides a relatively more mature platform for IoT applications (Amazon IoT) as compared to other vendors (though most of the approach described in this proposal is largely independent of this choice). Devices can connect to, and communicate with, applications running on the cloud over many different protocols. Device-specific SDKs are available for different languages. It offers support for reliable bidirectional messaging using the concept of device shadows to enable communication when the device network connectivity is not reliable. The cloud services that we are exploiting include (1) declarative, SQL-like rules engine to perform basic transformations of IoT data then reroute it to endpoints such as a storage container (S3 bucket), (2) support for triggers in the AWS event-based programming platform called Lambda, (3) redirect data streams into the Kinesis streams service to support real-time analytics, and (4) an archival storage service (Glacier).

### *C. Smart Intersections*

The raw video that is an input to our software is captured using fisheye cameras installed at traffic intersections. Compared to an ordinary video camera, a fisheye camera can capture the whole intersection in a wide panoramic and non-rectilinear image using its wide-angle fisheye lens. The camera transmits video to an edge-based GPU processor where image processing techniques convert the video to time-stamped 2D location coordinates of objects (vehicles, pedestrians, bicyclists).

Our video processing approach generates frame-by-frame detection and tracking of all the moving objects in an intersection. It also uses a temporal superpixel (supervoxel) method (Huang et al., 2018) to extract an accurate mask for object representations. These can be converted into trajectories that represent the spatial and temporal movement of traffic. A trajectory is a path traversed by a moving object that is represented as successive spatial coordinates and corresponding timestamps. Details of the video processing and analysis are provided in detail in Huang et al. (2019).

Using the output of the video processing, we have developed tools to process, filter, analyze, and display trajectories of vehicles and pedestrians passing through a traffic intersection. A trajectory of a moving object is its path represented by timestamped location coordinates of the object. For a typical, moderately busy intersection we studied, the volume of traffic is enormous, with over 10,000 trajectories being generated on a weekday. The number of trajectories quickly extrapolates to over a million trajectories per year for the intersection. For privacy protection, the information about the moving object is automatically anonymized by saving only the (relative) location coordinates of objects and, in the cases of vehicles, their size and color to our database.

Trajectories at intersections are dictated by the ongoing signaling status at an intersection. With the availability of advanced controllers that can record signal changes and detector events at a very high resolution (10 Hz), it is possible to generate a signal phase and timing log for the intersection for a given period.

The trajectories generated by video processing are uploaded in a MySQL database on the cloud. SPaT information is extracted from high-resolution controller logs and stored in a separate database as well.

The real-time trajectories, along with the current signaling state of the intersection, provide us valuable insights into any observed abnormal behaviors such as high propensities of signal light violations and risky maneuvers. If the intersection is large, two fisheye cameras may be installed to capture the complete intersection.

The trajectories are then clustered to derive normal and anomalous behavior. We compute the distance between two trajectories using a variation of FastDTW (Salvador and Chan, 2007) and a variation of the KMeans algorithm for clustering. The results of the clustering are then used by a visualization framework to display by phase, vehicle type, and time of day. An example is shown in Figure 2.

### *D. Smart network*

Traffic signals are controlled by sophisticated controller devices that collect a variety of data at every intersection.

**Figure 2:** Observed trajectories for different phases.

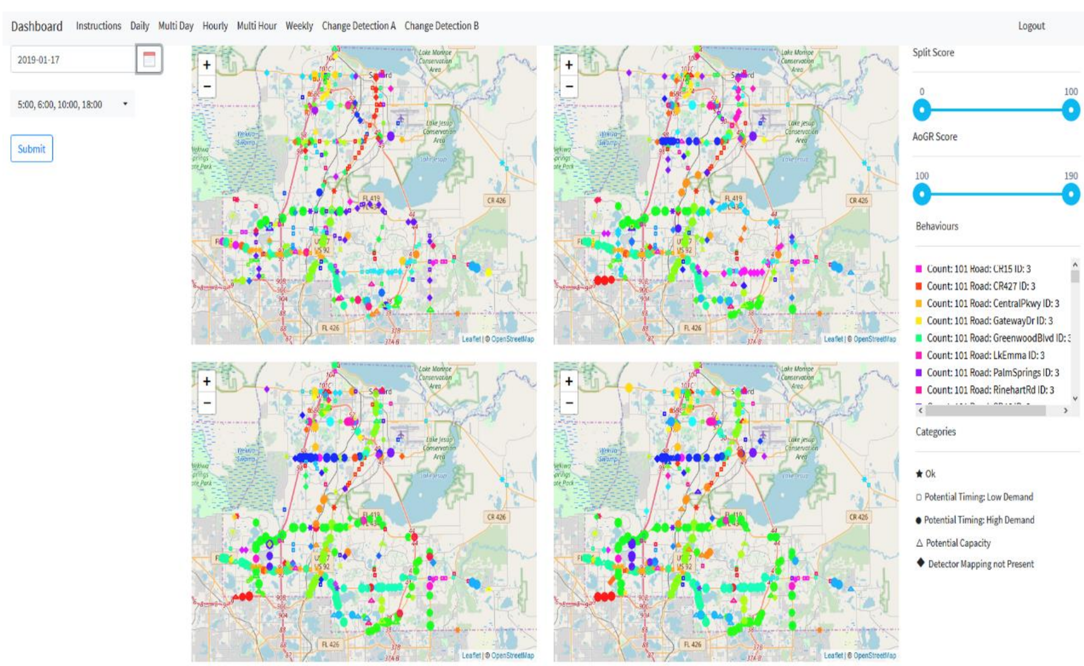


**Figure 3**: A dashboard highlighting the temporal recurrence of signal with similar behavior during the day. Specifically, the early morning/late night behavior can be contrasted with the daytime behavior.

Automated Traffic Signal Performance Measures (ATSPM), now being deployed in many traffic controllers, logs high-resolution (10Hz) data that opens a broad range of data analysis possibilities. We are using this data for a variety of applications.

We use demand-based split failures and arrivals on red as a Measure of Effectiveness (MOE) of an intersection. This can then be used for ranking or classifying signals. These measures can be used for determining level of traffic demand (based on split failures) and level of utilization of green time (based on the ratio of arrivals on red to arrivals on green).

We then apply clustering techniques to group together signals with similar MOes. Clustering is carried out in both space and time. Details are in Mahajan et. al. (2019). Figure 3 shows the dashboards of results on a single day. The intersections with similar performance are clustered together using the same color.

*Connected and Autonomous Vehicles (CAV)*

We have developed and deployed an intelligent controller for signalized intersections (Figure 4) at the FDOT Transportation Engineering Research Laboratory (TERL). The main components of the controller are an optimization algorithm for joint vehicle trajectory and signal control, a multi-sensor fusion system for obtaining real-time information from the surrounding traffic, and a signal controller for adapting the traffic signals based on the optimization results. The optimization algorithm considers both CAVs and conventional vehicles and produces signal timings that are field-implementable. Based on the real-time information from the sensor fusion component, it optimizes trajectories for connected and autonomous vehicles; the trajectories are transmitted to the vehicles with Dedicated Short-Range Communication (DSRC). The recommended trajectory is an ordered list of points and speeds that an autonomous vehicle can use to traverse the intersection, minimizing the system travel time. For connected vehicles, the message consists of speed recommendations for drivers.

Our simulation results (following work in Rosero, 2017, and Na, 2015) show that significant improvements can be achieved. For the tests at the isolated traffic intersection at TERL, we designed and implemented a multi-sensor fusion system based on DSRC and Doppler radar. We obtained the latitude, longitude, and speed information from vehicles equipped with DSRC. Using the radar, we were also able to get this information for conventional vehicles. These measurements were processed separately and converted into real-time tracks with a bank of Kalman filters and subsequently fused to produce accurate estimates of the state of each vehicle within range of the intersection. Our system (following Chavez-Garcia, 2016) was able to correctly classify all vehicles as automated, connected, or conventional during our tests, as well as estimate their position and speed with a high degree of confidence. We ran experiments with a vehicle equipped with a Cohda Wireless Mk5 on-board unit and a high precision GPS sensor to generate ground-truth data. We compared the tracking performance of different approaches configured with various vehicle kinematics models. Recently, we have started to experiment with vision-based multi-target tracking for traffic surveillance, with aims of developing a complete multi-sensor fusion system that combines video, radar, and DSRC. More details are in Li et. Al. (2014) and Emami et al. (2018).

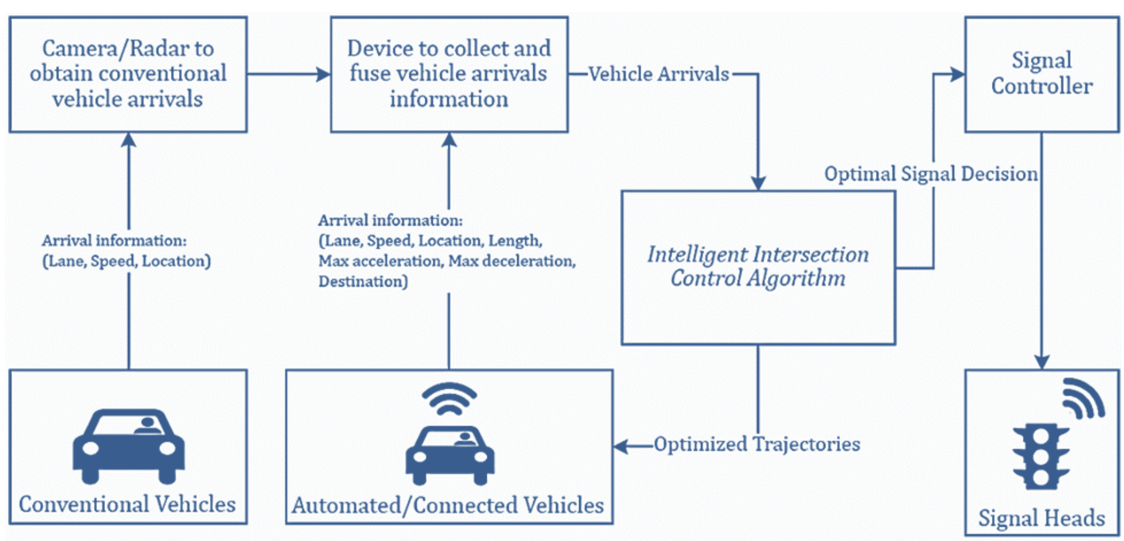


**Figure 4**: Overview of the signal control optimization algorithm with CAVs in the traffic stream.

*E. Security Issues*

We are using state of the art security mechanisms at the edge and the cloud. AWS provides cloud security (physical, filtering, and hypervisor security) while our systems are responsible for access control, encryption, etc.

## 7. Conclusions

Our work leverages the confluence of economical video sensors, significantly low-cost computing hardware, and cloud computing with open source analytics solutions to enable novel transportation applications. Hence, it will result in profound improvements in traffic management, smart city planning, and safety. The methods proposed will have a direct impact on video analysis for transportation problems and related disciplines and open avenues for research.

We will focus on emerging proactive measures for assessing and enhancing transportation safety and mobility. Regarding safety, unlike conventional methods which rely on crashes and, as such, represent a reactive approach to addressing safety problems, the use of space-time trajectories of vehicles and pedestrians to identify near-misses could help identify problems before actual crashes occur. While there has been significant recent interest in developing such trajectories, our ability to obtain and analyze continuous-time data at the network level will provide insights on how conflict points and patterns can change through the network. They can also determine changes over time even at the same intersection. This will allow us to identify dynamic safety improvement response strategies which were not possible using crash-based analytics.

The availability of continuous-time turning movements by vehicle type (including pedestrians) will also allow us to develop dynamic traffic operation strategies (signal timing plans) which can maximize system throughput while maintaining minimum desired performance levels at critical intersections. Finally, data on human-driven vehicle trajectories can provide baseline information to optimize trajectories of CAVs in the future.

This study goes well beyond visioning and algorithm developments into prototype development, deployment, and testing on a large-scale real-world system. The methods are sensitive to practical issues such as data quality and availability, technology constraints and interoperability, and institutional protocols. The effort also demonstrates the kinds of relationships that need to be developed among academicians, private sector vendors, and local and state agencies to ensure successful development and deployment of smart-city solutions.

## 8. Acknowledgements

This work is supported in part by the Florida Department of Transportation (FDOT), FDOT District 5, and NSF SCC 1922782. The opinions, findings, and conclusions expressed in this publication are those of the author(s) and not necessarily those of FDOT, the USDOT, or the National Science Foundation.

We are grateful to the City of Gainesville and District 5 for providing us access to a fisheye camera as well as for providing us with the video and controller logs at different intersections.

## References

[1] M. Abadi, A. Agarwal, et al., Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv preprint arXiv:1603.04467, 2016.

[2] R. O. Chavez-Garcia and O. Aycard. Multiple sensor fusion and classification for moving object detection and tracking. IEEE Transactions on Intelligent Transportation Systems. 17.2: 525- 534, 2016.

[3] P. Emami, M. Pourmehrab, M. Martin-Gasulla, S. Ranka, and L. Elefteriadou. A comparison of intelligent signalized intersection controllers under mixed traffic. ITSC. 341-348, 2018.

[4] X. Huang, C. Yang, S. Ranka, and A. Rangarajan. Supervoxel-based segmentation of 3D imagery with optical flow integration for spatiotemporal processing. IPSJ Transactions on Computer Vision and Applications. 10(1):1-6, 2018.

[5] X Huang, P. He, A. Rangarajan, and S. Ranka. Intelligent intersection: Two-stream convolutional networks for real-time near accident detection in traffic video. arXiv preprint arXiv:1901.01138, 2019.

[6] D. Mahajan, T. Banerjee, A. Rangarajan, N. Agarwal, J. Dilmore, E. Posadas, and S. Ranka. Analyzing Traffic Signal Performance Measures to automatically classify signalized intersections. VEHITS, 138-147, 2019.

[7] Z. Li, L. Elefteriadou, and S. Ranka, Signal control optimization for automated vehicles at isolated signalized intersections, Transportation Research Part C, vol. 49, pp. 1-18, December 2014.

[8] K. Na, J. Byun, M. Roh, and B. Seo. RoadPlot-DATMO: Moving object tracking and track fusion system using multiple sensors. Connected Vehicles and Expo (ICCVE), International Conference on. IEEE, 2015.

[9] M. Pourmehrab, L. Elefteriadou, and S. Ranka, Smart Intersection Control Algorithms for Automated Vehicles, Tenth International Conference on Contemporary Computing (IC3 2017), November 2017.

[10] L. A. Rosero and F. S. Osorio. Calibration and multi-sensor fusion for on-road obstacle detection. Robotics Symposium (LARS) and Brazilian Symposium on Robotics (SBR), Latin American. IEEE, 2017.

[11] S. Salvador and P. Chan. Toward accurate dynamic time warping in linear time and space. Intelligent Data Analysis. 11(5):561-80, 2007.