

Brief Industry Paper: The Necessity of Adaptive Data Fusion in Infrastructure-Augmented Autonomous Driving System

Shaoshan Liu^{*†}, Jianda Wang^{*‡}, Zhendong Wang[‡], Bo Yu[†], Wei Hu[†], Yahui Liu[†],
Jie Tang[§], Shuaiwen Leon Song[¶], Cong Liu[‡], Yang Hu^{‡||}
[†]PerceptIn, USA

[‡]The University of Texas at Dallas, USA

[§]South China University of Technology, China

[¶]University of Sydney, Australia

^{||}Corresponding author: Yang.Hu4@utdallas.edu

Abstract—This paper is the first to provide a thorough system design overview along with the fusion methods selection criteria of a real-world cooperative autonomous driving system, named *Infrastructure-Augmented Autonomous Driving* or IAAD. We present an in-depth introduction of the IAAD hardware and software on both road-side and vehicle-side computing/communication platforms. We extensively characterize the IAAD system in the context of real-world deployment scenarios and observe that the network condition fluctuates along the road is currently the main technical roadblock for cooperative autonomous driving. To address this challenge, we propose new fusion methods, dubbed "inter-frame fusion" and "planning fusion" to complement the current state-of-the-art "intra-frame fusion". We demonstrate that each fusion method has its own benefit and constraint. Adaptively choosing the fusion method according to the real-world condition will benefit the SoV without the violation of the SoV's safety requirements.

I. INTRODUCTION

The cooperative autonomous driving approach benefits from the collaboration between intelligent roads and smart vehicles enabled by emerging wireless communication technologies [1], [2]. It is not only more reliant but also more resilient in driving decision making compared to the traditional on-vehicle-only autonomous driving approaches, thanks to the driving-augmented information transmitted from the roadside systems. This paper provides the first thorough deployment experiences of an autonomous driving system enabled by intelligent roadside assistance, called *Infrastructure-Augmented Autonomous Driving (IAAD) system*.

The IAAD infrastructure consists of the roadside-deployed perception augmentation system, named *System-on-Road (SoR)*, and the on-vehicle driving automation system, named *System-on-Vehicle (SoV)*. The SoR aims to extend SoV's "vision" by transmitting the perception information to the SoV and it equips with the upgraded counterparts. By leveraging the integrated information from the SoR and the SoV, the IAAD system can make safer and more efficient driving decisions in real time than the on-vehicle-only autonomous system.

*Authors with equal contribution

Although the current IAAD design is promising, it is also very complex and comes with numerous real-world challenges. Our characterization study paints a frustrating picture that even the state-of-the-art version of the IAAD system configuration [3] is far below the standard for feasibility.

To the best of our knowledge, there is no existing work that focuses on the network latency and jitter affect on the performance of a real-deployed IAAD system. Since the SoV-SoR system is a cooperative system, a "synchronous process" is required for the SoV to wait for the data from the SoR before coming to the next pipeline module. The time when the SoR arrives at the SoV will be significantly affected by the network condition in the real-world. We observe that the real-world wireless network condition is complex and ever-changing. Specially, we find that the state-of-the-art fusion mechanism (intra-frame fusion) that fuses the SoR data to the SoV data with the same time stamp at the tracking module on the SoV cannot adequately handle the complex external factors including high network latency and jitter. Our deployment data indicates that the deadline miss ratio can be as high as 30%.

Based on our field experimental experience, we pinpoint that the system tolerance to the network impacts could be improved by carefully choosing the data fusion method. Therefore, to tackle the realtime network fluctuation challenge in IAAD design, we propose new fusion methods, dubbed "inter-frame fusion" and "planning fusion" to complement the current state-of-the-art "intra-frame fusion".

We demonstrate that each fusion method has its own benefit and constraint. Adaptively choosing the fusion method according to the real-world condition will benefit the SoV without the violation of the SoV's safety requirements.

II. REAL-WORLD IAAD TESTBED

The infrastructure-augmented autonomous driving (IAAD) system exhibits promising enhancement for the safety of the autonomous driving. *To the best of our knowledge, this is the first work that presents a real-world IAAD system and the corresponding thorough performance analysis.*

TABLE I: Hardware setup of IAAD.

	SoV	SoR
Industrial PC	CPU: Intel i9-9900K GPU: GTX 1060 x1 Memory: 32GB	CPU: Intel i9-9900K GPU: GTX 1080Ti x1 Memory: 32GB
Sensors	64-line LiDAR, camera x2, and Radar, GPS/IMU	Lightweight: 80-line LiDAR Heavyweight: 80-line LiDAR, camera x2, and Radars
Communication	OBU	RSU

In this section, we first describe the IAAD’s system hardware configuration and then introduce its algorithmic components and software pipeline in an end-to-end manner.

A. System Hardware Deployment

Hardware Setup: Table I lists the essential hardware configuration of our IAAD system. Our IAAD system consists of an on-vehicle driving automation system (System-on-Vehicle, SoV) and roadside-deployed driving augmentation system (System-on-Road, SoR). The SoV consists of an Industrial PC (IPC), a series of sensors, a communication component (on-board unit, OBU) as well as the ECU and the CAN bus connecting all components. ECU and CAN bus are standard autonomous components. The SoR consists of consists of an Industrial PC (IPC), a series of sensors, a communication component (roadside Unit, RSU). Different from the SoV, the SoR’s IPC is configured with a high-end Nvidia 1080Ti GPU and two optional sets of perception sensors. The SoR can be configured as heavy-weight SoR and light-weight SoR for the different deployment purposes. The heavy-weight SoR equips with cameras and LiDARs while the light-weight SoR only equips with LiDAR. This is because the LiDAR provides more accurate perception compared to the camera, especially during the night time. The RSUs and OBUs communicate directly to facilitate the fusion of the SoR data and the SoV data on the vehicle-side. The communication protocol is based on LTE-V2X currently and will support 5G-V2X in the future.

B. IAAD, An End-to-end Description

Fig. 1 presents an end-to-end illustration of our IAAD system. The main components’ latency and throughputs are demonstrated as well. Overall, both systems run a customized kinetic ROS middleware and a series of algorithmic modules. Specifically, our full-fledged L4 SoV includes Continental ARS 408-21 mmWave radars for detecting and tracking objects, cameras with 1920x1080 pixels for sensing and 2D perception module (based on YOLOv5), 64-line LiDAR and 3D perception module (based on PointPillars), GPS/IMU and localization module (based on NDT). The 2D perception and 3D perception results, as well as the outputs from mmWave radars, are combined to generate a comprehensive perception view in SoV-side sensor fusion module ①. The fused perception results are subsequently fed into the tracking module to keep track of each detected object ②. The purpose of tracking is to establish a temporal relation of the same object across multiple time frames so that the SoV gets to know the trajectory of the target object. Then the tracking results are fed to the prediction module, whose purpose is to predict the future trajectory of the target object based on its past trajectory ③.

Finally, the prediction results are fed to the planning module to make behavioral and motion planning decisions ④. After a planned trajectory is generated for SoV, the control module will work with the vehicle chassis to execute the trajectory ⑤.

In the meantime, the SoR detects and tracks objects with Continental ARS 408-21 mmWave radars. The SoR also performs 2D perception for the camera data and 3D perception for the LiDAR data. The SoR sensor fusion then generates the perception outputs ⑥. The perception outputs will be processed by tracking module and prediction module. Note that the SoR performs a heavy-weight prediction (H-prediction) [4]. Upon the various requests from SoV, either the perception outputs or prediction outputs of SoR could be transmitted to the SoV through LTE-V2X. These outputs can be fused with the SoV data and then fed to the corresponding modules on the SoV (e.g., either tracking ⑦ or planning ⑧).

III. THE REAL-WORLD NETWORK FLUCTUATION CHALLENGE

Our deployment experiences demonstrate that it is challenging to enjoy the benefit of long-range detection capability and high prediction accuracy in an IAAD system due to the real-world network fluctuation constraint. In this section, we thoroughly explore the challenges of an IAAD system provoked by the real-world network system constraint.

A. The New Metric Reflecting the Prediction Accuracy

To demonstrate the IAAD’s safety quantitatively, we firstly define a metric **displacement error**, which indicates the deviation between the predicted driving path and the ground-truth (i.e. ideal) driving path. The vehicle’s safety will be high when the displacement error keeps low. The equation 1 illustrates the definition of the displacement error:

$$Displacement\ error = \frac{\sum_{i=1}^N |Location(i)_{prediction} - Location(i)_{ideal}|}{N} \quad (1)$$

where the $Location(i)_{prediction}$ is the predicted location of the object i , the $Location(i)_{ideal}$ is the ground truth location of the object i and N is the total number of the detected objects. Our Field deployment results are illustrated in Table II. In this case, we assume every frame can arrive "in time" without the effect of the network fluctuation. We can observe that the displacement error decreases as the number of input frames increases. The displacement error will stay stable when the input frames achieve 20 frames. Consequently, our system will maintain a 20 frames window for the path prediction and will update the path prediction result when each new frame flows in.

TABLE II: Field deployment results

Input Frames Number	Displacement Error (meter)	Distance (meter) between vehicle and object		
		Without SoR	Light SoR	Heavy SoR
5	4.333	63	163	313
10	3.053	56	156	306
15	2.192	49	149	299
20	1.360	42	142	292
25	1.358	35	135	285

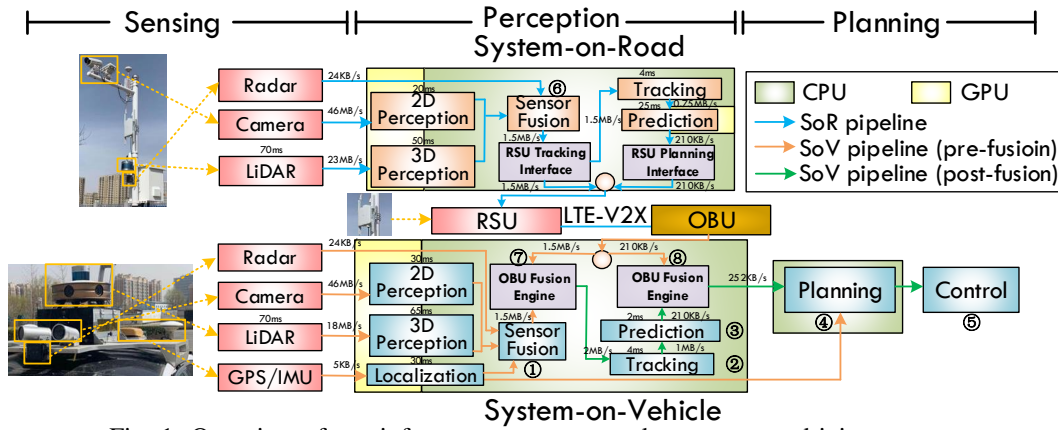


Fig. 1: Overview of our infrastructure-augmented autonomous driving system.

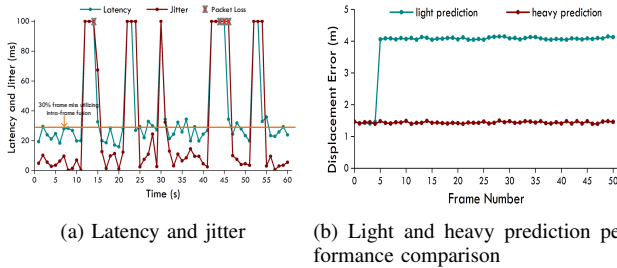


Fig. 2: Latency and jitter of real-world network and Comparison of light and heavy prediction

B. Real-World Network Fluctuation Challenge to the state-of-the-art Fusion Method

Intra-frame Tracking Fusion Constraint: The existing efforts on cooperative perception consider the tracking fusion as the state-of-the-art solution [3]. In the tracking fusion method, the RSU in SoR will extract the sensor fusion module’s output features and transmit them to SoV’s OBU using an LTE-V2X connection. After receiving the data, the OBU will fuse the data with its own sensor fusion module’s output and subsequently publish the fused data into the car-side software pipeline’s tracking module. Under this situation, the SoV output will be fused with the SoR frame which has the same timestamp. We name this tracking fusion method as “intra-frame tracking fusion”. Since the SoV-SoR system is a cooperative system, a “synchronous process” is required for the SoV to wait for the data from the SoR before coming to the next pipeline module. This synchronous process will require an extra time window on the SoV side to wait for the data from the SoR side. *The presumption of the existing work is that the frame from the SoR will always arrive the SoV in time. However, based on our field test, there are a bunch of cases when the SoV cannot receive the frame of the SoR in time due to the SoR latency caused by the network fluctuation.* Since the autonomous driving exists the End-to-End (E2E) pipeline boundary and the output interval boundary, the maximum time window has to comply with the above two boundaries to guarantee the vehicle’s safety. In the IAAD system, SoR latency caused by the network fluctuation will extend the E2E pipeline time and the output interval time.

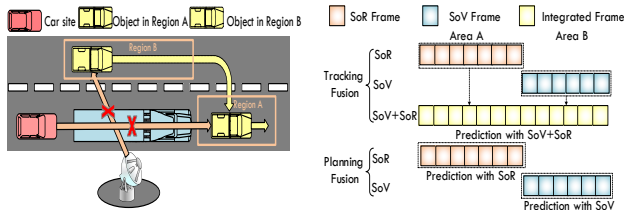
To avoid the violation of the two above boundaries, when the SoV cannot receive the SoR information within the time window, it needs to go through to the subsequent modules without the fusion to guarantee the safety of the vehicle. The E2E pipeline boundary and output interval boundary will be discussed as follows.

The E2E pipeline is defined as the duration from the time when the vehicle receives the sensing data to the time when the vehicle finishes its planning decision. Output interval is defined as the time between the two consecutive planning decisions. For the vehicle’s safety consideration, the E2E software pipeline latency and output interval both need upper boundaries. Thus, if the SoV cannot receive the data from the SoR “in time”, the data from SoV has to skip the “data fusion” process to go through the pipeline individually in order to follow the requirement of the E2E pipeline and output interval boundary.

Performance Degradation caused by the Real-world Network Fluctuation: Fig. 2a depict the network condition in our real test field. We can observe that the network latency fluctuates between 15 ms to 35 ms and sometimes will achieve 100ms. Even worse, in some cases the packets may totally get lost due to the bad network condition. Consequently, the SoR latency caused by the network fluctuation will probably lead to the violation of the E2E pipeline latency requirement. Besides, as shown in Fig. 2a, we observe that the network jitter will significantly hurt the stability of the vehicle’s planning decision interval time, which will cause the violation of the vehicle’s output interval boundary. Based on our field experience, the SoV will miss around 30% of the data from the SoR when simply utilizing intra-frame tracking fusion method, as illustrated in the Fig. 2a, which will cause the degradation of the performance of the IAAD system.

C. Alternative Fusion methods to tackle the network fluctuation challenges

In the previous section, we demonstrate the network fluctuation challenge confronted by utilizing the intra-frame tracking fusion method in the real world. Due to the E2E boundary and output interval boundary, this fusion method’s requirement is strict. Since the network will fluctuate in the real-world, a bunch of the frames will be missed from the SoR,



(a) Complementary view field case (b) Accuracy comparison
 Fig. 3: Analysis for fusion effectiveness under different fusion position

which will cause the performance degradation for our IAAD system. To deal with the network challenge and guarantee the IAAD safety, we take the step to explore a flexible data fusion framework, dubbed as *inter-frame tracking fusion and planning fusion*. In this section, we will illustrate the benefit and constraint of our inter-frame tracking fusion and planning fusion method compared to the intra-frame tracking fusion. **We demonstrate that the proposed fusion methods have their own benefit and constraint, adaptively choosing the fusion method according to the real-world condition will benefit the SoV without the violation of the SoV's safety requirements.**
Benefits and Constraints of Inter-frame Tracking Fusion: The intra-frame fusion incurs a strict waiting window for the SoV data. Based on our real deployment experience, the OBU in SoV can intermittently miss the waiting window due to some unavoidable network situations (e.g, base station handovers, congestion), which will cause the performance degradation in the system. To mitigate the affect, we explore the opportunity to utilize the previous SoR frame stored in the buffer to fuse with the current SoV frame if the intra-frame waiting window is missed. We name this mechanism as *inter-frame fusion*.

Unavoidably, the prediction accuracy of the inter-frame fusion will decrease compared to the intra-frame fusion (the requirement is most strict while the accuracy is the highest). However, if the timestamp of the frame from SoR (stored in the buffer) is very close to the current timestamp, the prediction accuracy does not decrease significantly. Besides, the tolerance of the prediction accuracy is high when concerning the far-fledged objects detected by the SoR compared to the nearby objects detected by the SoV. Consequently, the inter-frame fusion will effectively decrease the performance degradation compared to the system without fusion under the situation where the current frame from the SoR is missed.

However, the inter-frame fusion still has its constraint under the bad network condition, since the displacement error will become unacceptably high when missed frames accumulate due to the long-time network congestion. The tolerance time for the inter-frame fusion is around 0.6s, which will be adequate for the short time network congestion (i.e. base station handover) but still be not enough when long time network congestion occurs (i.e. go through tunnel region).

Benefits and Constraints of Planning Fusion: In our tracking fusion (intra-frame and inter-frame) mechanism, the post-perception data from the SoR will be combined with local data and processed by the tracking and *local prediction module*. However, a poor network environment can cause discontinued

provision of the SoR data due to the timeouts. The loss of augmentation of the SoR data will lead to a comparable displacement error in a much near distance with the object.

Since our local prediction module adopts a low-level prediction algorithm which only provides a prediction of next 0.5 seconds, the prediction accuracy will be significantly reduced if the network condition exacerbates, as shown in Fig. 2b. We can observe that the displacement error of light prediction can maintain in a low-level only for the first 5 frames and then increases to a high-level in the following 45 frames. To solve this issue, our SoR equips with a heavy-weight prediction algorithm ([4]) which can provide trajectory prediction for the next 5 seconds. This provides an opportunity that the prediction results from the SoR with lower displacement error, could be transmitted to the SoV and be utilized by planning module of SoV in the next 5 seconds instead of 0.5 second. As shown in Fig. 2b, the displacement error of heavy-prediction method keeps low in the entire 50 frames, and with a higher tolerance to the network latency. Motivated by this, we implement the planning fusion mechanism to overcome the performance degradation occurred in the tracking fusion method. By implementing the planning fusion, the tolerance time will be significantly extended compared to the tracking fusion.

Although the planning fusion will tolerate the congested network condition compared to the tracking fusion, the displacement error is higher than the tracking fusion. The reason for the higher displacement error is that the SoV's view and SoR's view are complementary with each other. Integrating SoR's data into a late stage in the SoV's software pipeline can cause the "spatial deficiency" for the system. For instance, as shown in Fig. 3a, when the object stays in region A, only the vehicle can detect the object (yellow vehicle) while the road cannot detect it. When the object goes into the region B, only the road can detect the object while the vehicle cannot detect it. As shown in Fig. 3b, the tracking fusion can integrate the fragmentary data from both sides to do its prediction while planning fusion will have to rely on the SoV and the SoR to do the prediction with their own fragmentary data, which will lead to the degradation of the accuracy for the final planning output. Consequently, when the network condition becomes better, we need to switch the planning fusion back to tracking fusion to avoid the spatial deficiency of the planning fusion.

REFERENCES

- [1] Shaoshan Liu, Bo Yu, Jie Tang, and Qi Zhu. Invited: Towards fully intelligent transportation through infrastructure-vehicle cooperative autonomous driving: Challenges and opportunities. In *Proceedings of the 58th Design Automation Conference*, 2021.
- [2] Shaoshan Liu, Liangkai Liu, Jie Tang, Bo Yu, Yifan Wang, and Weisong Shi. Edge computing for autonomous driving: Opportunities and challenges. *Proceedings of the IEEE*, 107(8):1697–1716, 2019.
- [3] Qi Chen, Xu Ma, Sihai Tang, Jingda Guo, Qing Yang, and Song Fu. F-cooper: feature based cooperative perception for autonomous vehicle edge computing system using 3d point clouds. In *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, pages 88–100, 2019.
- [4] Traffic predict. <https://gamma.umd.edu/researchdirections/autonomousdriving/traffic/>.