ICAT: An Indoor Connected and Autonomous Testbed for Vehicle Computing

Zhaofeng Tian¹, Yuankai He¹, Boyang Tian¹, Ren Zhong², Erfan Foorginejad², and Weisong Shi¹

Abstract— Indoor autonomous driving testbeds have emerged to complement expensive outdoor testbeds and virtual simulations, offering scalable and cost-effective solutions for research in navigation, traffic optimization, and swarm intelligence. However, they often lack the robust sensing and computing infrastructure for advanced research. Addressing these limitations, we introduce the Indoor Connected Autonomous Testbed (ICAT), a platform that not only tackles the unique challenges of indoor autonomous driving but also innovates vehicle computing and V2X communication. Moreover, ICAT leverages digital twins through CARLA and SUMO simulations, facilitating both centralized and decentralized autonomy deployments.

I. INTRODUCTION

The emergence of autonomous driving technologies has necessitated the development of specialized testbeds for their validation and improvement. Prominent outdoor autonomous driving testbeds like Mcity [1] at the University of Michigan, the edge computing testbed for autonomous driving at the University of Leeds [2], Singapore's Smart Mobility Testbed [3], and an autonomous shuttle testbed [4] built upon Taltech smart mobility city in Estonia, have been instrumental in facilitating this progress. These environments replicate realworld traffic scenarios, offering invaluable data and insights.

However, outdoor testbeds encounter several practical challenges that limit their efficacy and accessibility [5]. One of the most significant hurdles is the difficulty in creating realistic multi-agent traffic systems in a safe and controlled manner. Simulating dense, interactive traffic scenarios in outdoor settings poses substantial safety risks, making it challenging to test and refine autonomous systems under varied conditions. Additionally, the economic burden of establishing and maintaining such extensive outdoor infrastructure is substantial, often beyond the reach of smaller institutions and research entities.

Thanks to the more developed simulation technology, simulators like SUMO [6], CARLA [7], and LGSVL [8] generate high-fidelity autonomous driving testing scenarios for the community while circumventing the financial costs of building an outdoor testbed. The gap between real-world autonomous driving and simulation, however, still exists because of the inability of the simulation to run hardwarein-loop vehicle computing and V2X analysis. Besides the pure simulations, to tackle the efficiency dilemma in both financial and time costs of a real-world testbed, several indoor autonomous driving testbeds have been developed [9], [5], [10], [11], [12]. These testbeds are ingeniously downscaled to fit limited indoor spaces, simulating real-world scenarios. Utilizing motion-capturing systems, they provide centralized localization information directly from the scenario, bypassing decentralized sensor data processing. This streamlined approach significantly benefits research areas like path-trajectory planning, traffic optimization, and swarm intelligence. However, these testbeds are somewhat limited in their sensing and computing capabilities due to the absence of robust sensors and computing power, which leads to inefficiency in simulating real-world decentralized commuting-intensive autonomous driving tasks.

We introduce ICAT, an Indoor Connected and Autonomous Testbed (ICAT) in response to these limitations. ICAT stands apart by not only focusing on the unique challenges of indoor autonomous driving but also pioneering in vehicle computing and V2X communication. With digital twins, specifically leveraging the capabilities of CARLA and SUMO simulations, ICAT enables both centralized and decentralized autonomy deployments in simulated and realworld testbeds.

The merits of ICAT can be summarized into twomanifolds. (i) Connectivity and Vehicle Computing: Compared to other indoor autonomous driving testbeds, ICAT is emphasized on its V2X capability, which supports intervehicle, vehicle-infrastructure, and vehicle-server communications. While the on-board computing devices could serve as the resources for handling data from vehicles, infrastructure, and other information systems. (ii) Digital Twin and Multi-agent Simulation: The previous studies although built up a real-world multi-agent traffic or swarm system, faster developing process is prohibited without the full endorsement of a digital twin system. ICAT is not only integrated with CARLA, and SUMO interfaces but also a tailored pure Python multi-agent simulation environment, which releases the full potential in algorithm testing and iterating of autonomous driving tasks.

The rest of the paper is arranged as follows. The motivation and challenge of building ICAT are explained in Section II, related works are introduced and compared in Section III, and methods and technologies leveraged for designing ICAT are explained in Section IV. Two concrete case studies are conducted in Section V for validating the ICAT system. The design and implementation experience is then discussed in Section VI, and conclusions are drawn in Section VII.

 $^{^1} The CAR Lab, University of Delaware, Newark, USA {zhaofeng, willhe, tby, weisong}@udel.edu$

²Department of Computer Science, Wayne State University, Detroit, USA {ren, erfanf}@wayne.edu



Fig. 1. Panoptic view of ICAT platform in real world. ICAT is 6 by 5 meters in size, and 10 intelligent robots are used in ICAT now for autonomous driving studies. Traffic lights and signs are now integrated seamlessly with ICAT and enable interactions with traffic agents

II. MOTIVATION AND CHALLENGE

To better circumvent the unaffordable costs of building the outdoor autonomous driving testbed, and create more opportunities for the increasing interest in autonomy technologies, an indoor substitute would be an alternative to achieve the beginning stage of developing autonomous algorithms and modular structures. Moreover, thanks to the fast iteration of high-fidelity simulation platforms, developers could test their design in simulations conveniently. Although some indoor autonomous driving testbeds have been built, we have observed some challenges in previous studies building such a system, which are explained below.

(*i*) Localization Accuracy: An indoor testbed is highly downscaled compared to the real one, such that the localization accuracy in the downscaled road system plays a significant role in the downstream planning and control tasks. In some previous works, the localization is handled by a centralized motion-capturing vision system to extract the pose of each robot, such a method could limit the localization error to a smaller range of about 2cm. However, such a camera system for motion capture is very expensive, which brings more challenges for resource-limited entities. Using a decentralized method with sensors like 2D lidars to localize the robot is more affordable but introduces more challenges in accuracy.

(*ii*) Lack of On-board Computing Power: In the current studies, most robot agent running in the testbed does not have a proper computing device for decentralized autonomous driving, whereby the planning module would solely rely on a center computer or a server. This limits the research on vision and machine learning, which is inefficient in supporting decentralized information extraction and planning for each intelligent agent.

(*iii*) Digital Twin with Simulation Capability: The merit of a simulation platform for autonomous driving has been illustrated above. How to build a digital twin with access to such simulation software like CARLA and SUMO is vital for improving the quality and efficiency of the testbed. Unfortunately, the previous works do not support CARLA or SUMO interfaces.

To tackle these problems for a newer generation indoor autonomous driving testbed, we carefully choose different methods to solve them. The methods are explained in Section IV in detail.

III. RELATED WORK

In the domain of autonomous driving testbed, recent research has predominantly explored either simulation-based environments or real-world instances with limitations in scalability and decentralization. Studies like those in [13] and [14] have focused on single-vehicle testing. The former developed a system using RTK localization and AUTOSAR design on a remote-controlled car, effective for controller design testing but lacking in broader traffic system impacts. The latter introduced an autonomous delivery robot with a Gazebo environment, supporting real-world simulation collaborations, yet not addressing multi-agent autonomous driving capabilities.

On the pure traffic simulation front, AMoDeus, a testbed for autonomous transportation analysis is introduced in [15], which is confined to simulations without real-world applicability or sensory integration. Popular simulation platforms CARLA and SUMO enable 3D and 2D traffic generation and partial sensing capabilities. Both simulators have rich interfaces open to the public, however, they still provide simto-real transferring with limited resources.

Studies focusing on multi-robot systems, like [12] and [16], have made strides in swarm formation and autonomous driving systems. The former uses QR tags for localization but omits autonomous driving systems and onboard computing. The latter, despite including traffic signs and a digital twin map, is constrained by its design for cost-effectiveness, limiting independent localization and computing.

Research in indoor autonomous driving testbeds, such as in [5], [10], [9], [11] utilized motion-capturing systems for localization. Robots are controlled via a cloud server but are not equipped with independent driving capabilities in [5]. Study [10], while setting up a lane-free environment for multi-agent control algorithm testing, similarly did not enable decentralized methods. Scheffe et al. [9] introduces a scaled, remotely accessible autonomous driving environment. It supports multi-agent MPC-based collision avoidance in a centralized manner using a motion-capturing system. Nonetheless, it does not address decentralized autonomous

TABLE I								
TESTBEDS	BENCHMARK							

Research	Centralized	Decentralized	Independent Localization	V2X Communication	Vehicle Computing	Infrastructure	Digital Twin	Multi-agent	Carla / SUMO
Vedder et al. [13]	1	-	1	-	-	-	-	-	-
Tian et al. [14]	-	1	1	-	1	-	1	-	-
Ruch et al. [15]	1	-	-	-	-	-	-	1	-
Dosovitskiy et al. [7]	-	1	1	-	-	1	-	1	1
Krajzewicz et al. [6]	1	-	-	-	-	1	-	1	1
Pickem et al. [12]	1	-	-	-	-	-	-	1	-
Paull et al. [16]	1	-	-	-	-	1	1	1	-
Tran et al. [5]	1	-	-	-	-	-	-	1	-
Li et al. [10]	1	-	-	-	-	-	1	1	-
Scheffe et al. [9]	1	1	-	-	1	-	1	1	-
Stager et al. [11]	1	-	-	-	-	1	1	1	-
ICAT	1	1	1	1	1	1	1	1	1

driving with multi-sensors. Lastly, Stager et al.[11] describe a similarly scaled, remotely accessible autonomous driving environment, including simulated buildings. While it allows multi-agent centralized control with motion-capturing localization, the system is not designed for decentralized information processing and computing.

These works highlight a trend toward developing a scalable, decentralized autonomous driving testbed with onboard communication and computing capabilities. A brief comparison between our ICAT testbed and previous works is attached in Table I where independent localization shows that the localization relies on the onboard sensors, not a centralized motion-capturing system.

IV. ICAT DESIGN

In this section, different aspects of our proposed ICAT platform are introduced in detail. The methods and technologies that we have investigated for designing ICAT are illustrated as follows.

A. Digital Twin

Digital Twin technology is a rapidly evolving field with numerous motivations and applications across various industries. In the autonomous driving industry, road system planning, traffic management, and machine learning can be facilitated by digital twins where the detailed and accurate representation of physical assets, systems, or processes, allows for better analysis, simulation, and decision-making. This is particularly useful in complex systems where physical testing is impractical or costly.

The merits of leveraging a digital twin for our proposed ICAT platform can be summarized in three aspects. (i) Algorithm testing. Faster algorithm iteration can be achieved by efficient simulations. (ii) Traffic management. Managing a group of interacted traffic agents with a centralized autonomous driving algorithm in the real world hardly guarantees safety, the counterpart in a digital twin can release such safety burden. (iii) Remote testing. With advanced visualization capabilities, the information of testing in the real world can be accurately reflected on a remote end.

To build an indoor testbed for autonomous driving, a digital twin could support multi-purpose research and improve the developing efficiency. In such a digital design, however, we need to consider several problems. i) What road structure or features should be included? ii) What size is feasible to be implemented in the real world? iii) What interfaces it should have to integrate plug-ins of current autonomous driving simulators?

The proposed problems are tackled in sequence for our design. i) To meet the requirements of simulating intersections, merging points, and loops, we think a designed road system should include at least one intersection with leftturn, right-turn, and go-straight options for each direction. Also, merging points between lanes would support merging behavior research, and loops would simulate the traffic running in the system repeatedly with an automatic path planning algorithm. ii) We need to make sure the size of the real-world system can fit into the lab room. So finally the real-world system is sized at 6 by 5 meters, which is a 10-times downscale of a 60 by 50 digital system. Thereby, the finalized design is processed in RoadRunner [17], a road system building tool for autonomous driving development, and the final result is shown in Fig. 2(a). Correspondingly, we printed the digital map and deployed it in the lab, which is shown in Fig. 1. iii) To incorporate the convenience provided by current autonomous driving simulators, we build simulations using the designed ICAT map in both CARLA and SUMO, such that the built-in functions and algorithms can be leveraged by ICAT in a plug-and-play way. The CARLA and SUMO simulation demonstrations are shown in 2(b)(c), where the traffic could be either controlled by a traffic management algorithm from CARLA or SUMO.

To support user-defined simulation scenarios, we have also developed a pure Python simulation system for ICAT, which enables the development of traffic management, trajectory planning, and control algorithms. One of the use cases is demonstrated in Fig. 3, where a traffic management system with a trajectory planning algorithm is developed, and this process will be further illustrated in Section IV-D

B. Infrastructure and V2X

ICAT also contains a down-scaled infrastructure system. To ensure a realistic representation, multiple traffic signs and traffic lights are used as Fig. 1 shown. The traffic signs include speed limit, yield sign, construction zone, stop sign,



(a) ICAT road system designed in RoadRunner

(b) ICAT simulation in CARLA

(c) ICAT simulation in SUMO

Fig. 2. Digital twin system of ICAT. (a) ICAT's map is designed within RoadRunner [17], which is a 3D interactive autonomous driving system builder. (b) The designed map then can be seamlessly imported into CARLA simulations. (c) SUMO simulations are also enabled with the ICAT road system.

and do not enter signs.

The infrastructure information is also shared following the J2735 C-V2X standards. To accurately mimic realistic V2X communication, ICAT leverages the in-door WiFi system for communication. WiFi systems are similar to outdoor C-V2X systems due to their broadcasting nature. An underlying ROS communications stack is fitted into the WiFi system. A network model is used to mimic C-V2X message latency and packet loss.

Map information such as lane and intersection location and allow maneuvers are broadcasted by the infrastructure nodes. We've also added a Raspberry Pi controller to each traffic light signal so that traffic light state and timing information can be broadcast to nearby test vehicles. Furthermore, the J2735 standards also allow emergency vehicles to change traffic light states. This can be replicated with the help of the WiFi network, ROS communications stack, and the Raspberry Pi.

C. Localization

Compared to using an expensive centralized motion capturing system for localization, based on the onboard 2D Lidar, an ICAT robot can conduct localization using some mature algorithm in an efficient way. Additionally, a decentralized localization method could endorse richer testing scenarios for autonomous driving offering more realistic sensing inputs.

Specifically, to enable safe navigation of ICAT, we need to establish a solid localization foundation, including two functions, flexible global re-localization and real-time accurate pose tracking to produce reliable pose information of vehicle status. Global re-localization refers to obtaining the rough pose in the map with only the most recent observation when the localization system of ICAT is just booted or the tracking function fails. After producing an initial pose, the pose tracking function continues to estimate the movement of sensors periodically to ultimately calculate the vehicle status such as speed and orientation relative to the map.

Currently, there are many SLAM systems supporting both localization functions. For example, the ORB SLAM family

[18], [19] which is a group of camera-based SLAM systems can locate the 3D world position of ORB features from images to track camera movements. By recording historical features into a library of historical pose, the ORB SLAM can achieve global re-localization by searching most similar features. However, this solution is less effective in an indoor environment as shown in a common indoor room as shown in the figure. Due to the limited performance of the camera depth sensor, the system found it hard to recognize the same wall at different distances. As a result, ORB SLAM is easy to commit tracking failure. Compared to camera-based SLAM, LiDAR-based SLAM [20], [21] is better at pose tracking since the sensor data includes more accurate and wider depth information, but is worse at global re-localization due to the challenge of pose description especially for 2D LiDAR equipped by ICAT whose data is sparse and lack of geometrical structure.

In our implementation of the localization framework, we use Normal Distribution Transform (NDT) SLAM [22] which is designed for 3D LiDAR originally to build a tracking layer due to its reliable performance compared to traditional 2D LiDAR SLAM such as [23], [24]. Furthermore, the generated layer is a 2D point cloud map that can be directly re-used by all robots in an ICAT fleet. NDT is also responsible for obtaining tracking results after obtaining the initial pose from global re-localization which is based on ORB features and historical pose library in the format of Bags of Words [25]. Inspired by ORB SLAM, global relocalization will first extract ORB features from the latest image to calculate a visual word for searching the nearest word in the historical pose library. The found image is then matched with current features to estimate a rough pose. NDT SLAM can further optimize this pose to promise later tracking quality.

D. Centralized Traffic Management System

CARLA and SUMO are designed primarily for autonomous vehicle research and development. Its traffic management system is created to simulate realistic traffic scenar-



Fig. 3. ICAT traffic management system testing in a Python simulation, where simulated 10 traffic agents could be controlled by a centralized traffic management algorithm. The planned trajectories for different robots are marked in different colors.

ios, providing a safe, controllable, and flexible environment for testing autonomous driving systems and algorithms.

Both CARLA and SUMO primarily operate with a centralized architecture. In this system, the simulation environment, including traffic scenarios, vehicle behaviors, and environmental conditions, is controlled by a central server or application. The control over individual vehicles, pedestrians, and other elements can be scripted or managed through the central API, allowing for precise control and manipulation of the simulation environment. Specifically, in the context of autonomous vehicle testing, this centralized approach allows for a centralized planning framework and cooperative collision avoidance.

The merit of using a centralized traffic management system can be summarized as follows. i) Capable of simulating large networks of roads and junctions, complete with public transport, pedestrians, and varying traffic densities [26]. ii) Supports traffic light control, route planning, and vehicle-toinfrastructure communication. iii) Can better support traffic flow analysis, such as measuring travel times, system throughout, and emissions [27].

However, in the implementation of CARLA's traffic management system, we found a deadlock problem when the number of cars increased, and traffic efficiency degraded significantly. The reason is that in the intersections or merging points, the checking points are preset such that the first car hitting the checking point will go first leaving the other car totally stopped to yield. Such a First-In-First-Out(FIFO) [28] policy is not optimal, and further optimization is needed.

In ICAT, we build such a traffic management system independently from CARLA and SUMO but refer to some ideas from them, i.e., the node-based traffic system graph. This graph will determine a full junction system of the roads, and support graph-based path search. Moreover, to tackle the problems of deadlocks and low efficiency, we leverage multiagent trajectory planning to dynamically tune the speed for intersection/merging cars to pass at a higher efficiency.

The process of operating such a traffic manager system is described in Algorithm 1. First, A traffic system graph is built up with nodes and edges, where each node is connected with others with a directed edge either shaped as a curve or

Algorithm 1 Optimized Traffic Management System

- 1: Initialize traffic graph G, car states X, path buffer P, trajectory buffer Traj.
- 2: Initialize traffic signals, obstacles states O.
- 3: for each car i do
- 4: **if** P[i] is empty or invalid **then**
- 5: $P[i] \leftarrow A^*(G)$
- 6: **end if**
- 7: $(s_i, d_i) \leftarrow \text{LOCALIZETOSD}(P[i])$
- 8: $\operatorname{Traj}[i] \leftarrow \operatorname{POLYNOMIALPLANNING}(s_i, d_i, X_i, O)$
- 9: conflicts \leftarrow COLLISIONDETECTION(Traj[i], Traj)
- 10: **for** each conflict in conflicts **do**
- 11: $\operatorname{Traj}[i] \leftarrow \operatorname{REPLANTRAJECTORY}(i, \operatorname{conflict})$
- 12: end for
- 13: end for
- 14: Update global state based on new trajectories and environment changes for the next time step.



Fig. 4. Graph topology for ICAT traffic system, consisting of nodes and edges that a pair of nodes is connected with either a straight line or a curve. Waypoints are discretized with a fixed distance as an attribute of the edge.

straight line shown in Fig. 4. Then A* search can be applied to each initialized robot to automatically generate paths and update goals (when approaching the previous goals) such that the simulation of a traffic system would not stop at reaching goals.

After that, the coordinates of (x, y) in Euclidian coordinates need to be transformed into (s, d) Frenet coordinates [29], such that the location of cars can be associated with roads in a convenient way. Based on the Frenet localization, the trajectory for each car is planned with polynomial interpolation between the current state and the future state. Other cars' states are also considered for trajectory planning, for example, the leading car's location and speed will be taken into account for doing an Adaptive Cruise Control (ACC) style trajectory planning. Other than ACC, the logic of trajectory planning for merging and diverging cases will be considered independently. A visualization of the planned trajectories is shown in Fig. 3

Finally, the conflicts among planned trajectories will be addressed using a replanning process, where a non-linear optimization problem is usually set up to minimize the sum of the inter-car distance between conflicted cars for collision avoidance [30]. The implementation of our designed traffic management system is illustrated in Section V.

E. Decentralized Autonomous Driving

Decentralized Autonomous Driving differs from traffic management methods in that it doesn't rely on externally received command information, the traffic agents solely rely on their own sensors and computing devices to conduct perception and data processing. Thanks to heterogeneous on-board computing devices, the ICAT robots with ARM computing platforms leverage Scalable Open Architecture For the Embedded Edge (SOAFEE) [31] and Autoware [32] OpenAD Kit to support autonomous driving. SOAFEE and OpenAD Kit serve as a baseline for autonomous driving and enable researchers to quickly and easily modify the perception, planning, and control algorithms. Robot vehicles with other computing architectures leverage Autoware.AI, an autonomous driving stack based on ROS1, as the baseline.

Besides the mature autonomous driving pipelines like Autoware and SOAFEE, ICAT robots incorporate advanced vision reasoning technology for object detection and endto-end navigation. Similar to Tesla's current solution, the current decentralized method uses depth cameras for environmental data perception [33]. To ensure smooth driving along the road boundaries, we employ Color-based Lane Keeping [34]. The process begins with calibration, addressing distortions in the camera image caused by radial and tangential deformations. Calibration utilizes a known spatial relationship (calibration board) to deduce the camera's intrinsic parameters, correcting distortions in raw images.

Following calibration, we apply distortion correction to raw images and implement real-time HSV control for color tracking. Adjusting HSV thresholds filters out interfering colors, allowing for ideal recognition of squares in complex environments. Color transforms and gradients create a thresholded binary image.

Additionally, due to the limited variety and quantity of traffic sign shapes, we employ a lightweight YOLOv5 [35] network for single-stage object detection to recognize these signs. We perform data augmentation by cropping standard traffic signs and introducing slight variations in brightness, angles, and horizontal flips, while simultaneously generating labeled bounding boxes. The lightweight YOLOv5 network helps reduce the computational load on the experimental vehicle, decrease space requirements, lower overall system energy consumption, and achieve real-time output speeds.



Fig. 5. Our designed HydraT platform, is equipped with an advanced Lidarcamera sensing system and empowered with up-to-date navigation software.

F. Vehicle Computing and Multi-user Management

Besides the intelligent transportation and autonomous systems highlighted in the ICAT platform, the ICAT platform implements an innovative integration of vehicle computing and multi-user management. As the vehicle computing power is increasing rapidly, vehicle can serve as mobile computing platforms for various tasks. [36]. Adapting such an idea, our platform harnesses a fleet of 10 sophisticated robots with heterogeneous computing devices including Nvidia Nano, TX2, and NX. Each robot serves a dual purpose for computing: participating in a dynamic traffic system and acting as a resource for distributed computing tasks, such as federated Machine Learning (ML) [37].

At the core of our platform lies the autonomous driving algorithms, which are rigorously tested within a simulated traffic system. For example, in decentralized autonomous driving, perception, and planning tasks will be processed onboard, where decentralized computational power plays an important role for autonomous agents to make decisions in real-time, adapting to the continuous flow of the traffic system.

Beyond traffic management, our robots with advanced computational capabilities, allowing them to engage in federated ML tasks during their "commute". This innovative approach utilizes the otherwise idle processing power, turning each robot into a node within a distributed learning network. By doing so, we enhance the utility of our robotic fleet, contributing to a range of computational tasks for multiuser cases, from optimizing traffic algorithms to processing large-scale data analyses, all while ensuring the integrity and privacy of the data involved.

To address the complexities of multi-user management, our platform leverages Docker containers, providing a robust and isolated environment for each user. These containers encapsulate the necessary tools, libraries, and runtime for users to deploy and manage their autonomous driving and ML applications without interference. This separation is crucial for maintaining system stability and security, ensuring that the computational activities of one user do not adversely affect those of another.

G. HydraT Platform and Auto-recharge

Inspired by [38], [14], we developed a scaled autonomous car platform, Hydra-T, suitable for experiments in the ICAT environment shown as Fig. 5 While having the same Ackermann-steering system as a real car. To facilitate accurate sensing, the environment information can be extracted from both an accurate industry-level 2D-LiDAR, and an Intel RealSence depth camera. Combined with a 9-DOF IMU, the accuracy of localization and pose estimation can be further enhanced. Supporting intensive vehicle computation for multi-purposes, the robot is designed to adapt an Nvidia Orin machine that provides a powerful online machine inference capability.

To reduce maintenance requirements we design a custom charging dock for the robot, such that the robot can conduct automatic recharge before the battery is running out. This eliminates the need for researchers' intervention since the charging feature enables the robot to dock and recharge itself autonomously similar to vacuuming robots, thus achieving a higher time efficiency for a multi-agent system. Furthermore, the charging stations can also be a part of experiments simulating the real-world problems of electric vehicles with charging path planning and optimization for energy management.

V. CASE STUDIES

The ICAT platform's efficacy is evaluated by two case studies: the operational integrity of the traffic management system and the execution of federated machine learning (ML) algorithms. These case studies are designed to validate the ICAT platform's functionality and showcase its versatility in handling complex computational tasks while managing a fleet of autonomous robots.

A. Traffic Management System

As we explained in Section IV-D, the whole management system is first built in a pure Python environment for validation. As shown in Fig. 3, running 10 robots simultaneously in ICAT allows each car to automatically plan a path with random start and stop nodes with an A* algorithm. Then the car plans trajectories that dynamically adjust space and speed to a leading car in an ACC logic. The planned trajectory is rendered as apparent frames in the front of cars.

However, in real-world experiments, the communications between cars and the centralized server are implemented in ROS [39] system, which is different in simulations. To bridge this transferring gap between simulations and ROS real-world tests, we first assign the different robot namespace from *robot1* to *robotn* based on the number of robots used, hereby all the robot topics will be separated under their unique namespace. Next, different functions are separated into different ROS nodes, for instance, the localization module and trajectory planning module will publish poses and trajectories from different nodes.

Besides interface issues, real-world validations, localization, and trajectory following (control) problems need to be handled carefully. This is because, in the simulation, the position/pose information is captured from the defined vehicle dynamic models without any sensors and control noises, which is assumed as ground truth. Assuming perfect trajectory tracking for every robot enables the robot to update its state perfectly to follow the trajectory, which is impossible in the real world where each frame in the trajectory necessitates accurate following by a well-designed control module.

To respectively conquer the localization and control problem, we leveraged the NDT method [22] for localization and observed an accuracy improvement from the original AMCL method [40], which has been explained in Section IV-C. For the control method, we have implemented a pure pursuit method [41] to track the planned trajectory, where the controller takes a look-forward distance as input to chase the look-forward point, and hence improve the robustness against localization noises by ignoring the very close errors.

To better visualize the real-world experiments and support monitoring at the digital end shown in Fig. 6, we developed an RViz-based monitoring system where each car's trajectory and ICAT road system will be published as topics to be shown on screen. Furthermore, a pose initialization tool is integrated into the RViz, such that an initial pose for localization can be easily achieved by pulling the mouse toward the corresponding pose in the real world.

B. Federated Machine Learning

The second aspect of our evaluation focused on the execution of federated ML tasks. Each robot, during its simulated commute, participated in a distributed ML training session. The task involved collaboratively learning a predictive model while keeping on each robot the training data which is provided by our industry partner.

The experimental evaluation of the ICAT platform confirmed its robustness and versatility. The proposed traffic management system proved efficient, safe, and scalable. Simultaneously, the platform's capability to perform federated ML tasks without compromising its primary functions was successfully demonstrated. These results underscore the potential of ICAT as a dual-purpose platform, catering to the needs of intelligent traffic management and vehicle computing research.

VI. DISCUSSION AND FUTURE WORK

This paper presents the innovative ICAT platform and details our extensive experimentation and setup processes. Through our testing and evaluation, we have gained valuable insights into the complexities of autonomous vehicle technology. Our experiences and findings can be summarized in the following key observations:



(a) Running a 4-robots traffic management demo

(b) 4 robots in RViz

(c) 6 robots in RViz

Fig. 6. An illustration of traffic management real-world experiments. (a) 4 robots are running a traffic management demo. (b) Real-world demos can be also monitored in RViz, where the ICAT road is published as a Path message (the tilted lines are caused by the automatic connection of section start and end points in the Path message), the planned car trajectories from the central manager are rendered as different colors. (c) We have tested up to 6 robots in the ICAT real-world traffic system.

(i) Impact of NDT Localization Noise on Pose Initialization. One of the significant challenges encountered during our experiments was the influence of noise in NDT localization on pose initialization. We observed that even minor noise in NDT measurements could lead to inaccuracies in initial pose estimation. This, in addition, also affected the navigation accuracy of the autonomous vehicles on our platform. The findings underscore the need for more robust noise filtering techniques in NDT localization to enhance the system precision.

(*ii*) Control Challenges in Trajectory Tracking. Another key observation was related to the control aspects, particularly in tracking the planned trajectory using a pure-pursuit controller. While the pure-pursuit algorithm is a widely accepted method for trajectory following in autonomous vehicles, we noticed that the controller still exhibited errors in trajectory tracking, particularly on a curvy path.

(*iii*) Effect of Communication Lag on Response Speed. Lastly, our experiments brought to light the significant impact of communication lag on the response speed of the cars. The communication delay can have critical implications in a system where multiple autonomous vehicles communicate and operate synchronously. We observed that even minimal lags could lead to noticeable delays in the vehicles' responses to control commands. This finding is particularly pertinent in scenarios requiring fast decision-making and rapid maneuvering.

In our future works, the insights we have obtained will be leveraged to improve the performance of the ICAT platform. Our future research plans will include the following items. First, optimize the localization accuracy with better filtering tricks. Second, a localization error-considered model predictive control method will be implemented to minimize the trajectory tracking errors. A learning method will acquire an accurate spatial-temporal environment dynamic model based on the collected real-time data. Finally, we will equip ICAT with more powerful communication devices, increasing the bandwidth at the hardware end. Furthermore, task scheduling and lag-considered safety protection will be investigated to avoid the harmful impact of communication delays.

VII. CONCLUSION

This paper presented the ICAT platform, an advanced indoor testbed developed for cutting-edge research in autonomous vehicle systems and advanced vehicle computing. Through our evaluations, ICAT has demonstrated its proficiency in managing simulated traffic systems and executing complex federated ML tasks, proving its value as a versatile tool for the intricate demands of modern intelligent transportation and computational research.

A key strength of the ICAT platform lies in its integration of connectivity, vehicle computing, digital twin technology, and multi-agent simulation. The platform's exceptional V2X (Vehicle-to-Everything) capabilities and advanced onboard computing devices allow ICAT to adeptly manage and process data from various sources, including inter-vehicle, vehicle-infrastructure, and vehicle-server communications.

Furthermore, ICAT's digital twin system, integrated with CARLA and SUMO interfaces and a tailored pure Python multi-agent simulation environment, sets it apart from previous studies. This integration allows for the simulation of realworld multi-agent traffic or swarm systems and significantly accelerates the development and iteration of algorithms for autonomous driving tasks.

As we look to the future, we will focus on expanding ICAT's capacity, enhancing its simulation fidelity, and integrating decentralized computing paradigms more deeply. These advancements are directed towards supporting the evolving field of autonomous vehicle research, ensuring that ICAT remains a pivotal and invaluable resource for researchers and practitioners in this rapidly advancing area.

ACKNOWLEDGMENT

The work is partially supported by the US National Science Foundation (NSF) under Grant No. 2311087. Special thanks are extended to Toyota and Tier IV for their support. The authors also wish to express their gratitude to all individuals and organizations that contributed directly or indirectly to the success of this work.

REFERENCES

- Y. Dong, Y. Zhong, W. Yu, M. Zhu, P. Lu, Y. Fang, J. Hong, and H. Peng, "Mcity data collection for automated vehicles study," *arXiv* preprint arXiv:1912.06258, 2019.
- [2] Z. Nezami, E. Pournaras, A. Borzouie, and J. Xu, "Smotec: An edge computing testbed for adaptive smart mobility experimentation," *arXiv* preprint arXiv:2307.11181, 2023.
- [3] "Ntu smart mobility testbed," https://www.ntu.edu.sg/ciss/ research-capabilities/research-programmes/cosmo/research/wp1, accessed: 2023-09-30.
- [4] R. Sell, R.-M. Soe, R. Wang, and A. Rassõlkin, "Autonomous vehicle shuttle in smart city testbed," in *Intelligent System Solutions for Auto Mobility and Beyond: Advanced Microsystems for Automotive Applications 2020.* Springer, 2021, pp. 143–157.
- [5] D. Tran, E. Tadesse, P. Batapati, W. Sheng, and L. Liu, "A cloud based testbed for research and education in intelligent transportation system," in 2015 IEEE International Conference on Robotics and Biomimetics (ROBIO). IEEE, 2015, pp. 452–457.
- [6] D. Krajzewicz, G. Hertkorn, C. Rössel, and P. Wagner, "Sumo (simulation of urban mobility)-an open-source traffic simulation," in *Proceedings of the 4th middle East Symposium on Simulation and Modelling (MESM20002)*, 2002, pp. 183–187.
- [7] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.
- [8] G. Rong, B. H. Shin, H. Tabatabaee, Q. Lu, S. Lemke, M. Možeiko, E. Boise, G. Uhm, M. Gerow, S. Mehta *et al.*, "Lgsvl simulator: A high fidelity simulator for autonomous driving," in 2020 IEEE 23rd International conference on intelligent transportation systems (ITSC). IEEE, 2020, pp. 1–6.
- [9] P. Scheffe and B. Alrifaee, "A scaled experiment platform to study interactions between humans and cavs," in 2023 IEEE Intelligent Vehicles Symposium (IV). IEEE, 2023, pp. 1–6.
- [10] B. Li, D. Cao, S. Tang, T. Zhang, H. Dong, Y. Wang, and F.-Y. Wang, "Sharing traffic priorities via cyber–physical–social intelligence: A lane-free autonomous intersection management method in metaverse," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 53, no. 4, pp. 2025–2036, 2022.
- [11] A. Stager, L. Bhan, A. Malikopoulos, and L. Zhao, "A scaled smart city for experimental validation of connected and automated vehicles," *IFAC-PapersOnLine*, vol. 51, no. 9, pp. 130–135, 2018.
- [12] D. Pickem, P. Glotfelter, L. Wang, M. Mote, A. Ames, E. Feron, and M. Egerstedt, "The robotarium: A remotely accessible swarm robotics research testbed," in 2017 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2017, pp. 1699–1706.
- [13] B. Vedder, J. Vinter, and M. Jonsson, "A low-cost model vehicle testbed with accurate positioning for autonomous driving," *Journal* of Robotics, vol. 2018, pp. 1–10, 2018.
- [14] Z. Tian and W. Shi, "Design and implement an enhanced simulator for autonomous delivery robot," in 2022 Fifth International Conference on Connected and Autonomous Driving (MetroCAD), 2022, pp. 21–29.
- [15] C. Ruch, S. Hörl, and E. Frazzoli, "Amodeus, a simulation-based testbed for autonomous mobility-on-demand systems," in 2018 21st International Conference on Intelligent Transportation Systems (ITSC). IEEE, 2018, pp. 3639–3644.
- [16] L. Paull, J. Tani, H. Ahn, J. Alonso-Mora, L. Carlone, M. Cap, Y. F. Chen, C. Choi, J. Dusek, Y. Fang *et al.*, "Duckietown: an open, inexpensive and flexible platform for autonomy education and research," in 2017 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2017, pp. 1497–1504.
- [17] MathWorks, "Roadrunner," https://www.mathworks.com/products/ roadrunner.html, 2023.
- [18] R. Mur-Artal, J. Montiel, and J. Tardos, "Orb-slam: a versatile and accurate monocular slam system," *IEEE Transactions on Robotics*, vol. 31, pp. 1147 – 1163, 10 2015.
- [19] R. Mur-Artal and J. Tardos, "Orb-slam2: an open-source slam system for monocular, stereo and rgb-d cameras," *IEEE Transactions on Robotics*, vol. PP, 10 2016.
- [20] G. Kim, S. Yun, J. Kim, and A. Kim, "Sc-lidar-slam: A front-end agnostic versatile lidar slam system," 02 2022, pp. 1–6.

- [21] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "Fast-lio2: Fast direct lidar-inertial odometry," *IEEE Transactions on Robotics*, vol. 38, pp. 1–21, 08 2022.
- [22] M. Magnusson, "The three-dimensional normal-distributions transform — an efficient representation for registration, surface analysis, and loop detection," Ph.D. dissertation, 12 2009.
- [23] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with rao-blackwellized particle filters," *Robotics, IEEE Transactions on*, vol. 23, pp. 34 – 46, 03 2007.
- [24] S. Kohlbrecher, J. Meyer, O. von Stryk, and U. Klingauf, "A flexible and scalable slam system with full 3d motion estimation," in *Proc. IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR).* IEEE, November 2011.
- [25] D. Galvez-Lopez and J. Tardos, "Bags of binary words for fast place recognition in image sequences," *Robotics, IEEE Transactions on*, vol. 28, pp. 1188–1197, 10 2012.
- [26] B. Chalaki and A. A. Malikopoulos, "Optimal control of connected and automated vehicles at multiple adjacent intersections," *IEEE Transactions on Control Systems Technology*, vol. 30, no. 3, pp. 972– 984, 2021.
- [27] J. Rios-Torres and A. A. Malikopoulos, "A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 5, pp. 1066–1077, 2016.
- [28] A. Mirheli, M. Tajalli, L. Hajibabai, and A. Hajbabaie, "A consensusbased distributed trajectory control in a signal-free intersection," *Transportation research part C: emerging technologies*, vol. 100, pp. 161–176, 2019.
- [29] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, "Optimal trajectory generation for dynamic street scenarios in a frenet frame," in 2010 IEEE international conference on robotics and automation. IEEE, 2010, pp. 987–993.
- [30] M. A. S. Kamal, J.-i. Imura, T. Hayakawa, A. Ohata, and K. Aihara, "A vehicle-intersection coordination scheme for smooth flows of traffic without using traffic lights," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 3, pp. 1136–1147, 2014.
- [31] SOAFEE, "Software Defined Vehicle," https://www.soafee.io/blog/ 2022/sdv_with_cloud, 2023, [Online; accessed 23-Dec-2023].
- [32] S. Kato, S. Tokunaga, Y. Maruyama, S. Maeda, M. Hirabayashi, Y. Kitsukawa, A. Monrroy, T. Ando, Y. Fujii, and T. Azumi, "Autoware on board: Enabling autonomous vehicles with embedded systems," in 2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS). IEEE, 2018, pp. 287–296.
- [33] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Y. Qiao, and J. Dai, "Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers," in *European conference on computer vision*. Springer, 2022, pp. 1–18.
- [34] A. B. Hillel, R. Lerner, D. Levi, and G. Raz, "Recent progress in road and lane detection: a survey," *Machine vision and applications*, vol. 25, no. 3, pp. 727–745, 2014.
- [35] Ultralytics, "YOLOv5: A state-of-the-art real-time object detection system," https://docs.ultralytics.com, 2021.
- [36] S. Lu and W. Shi, "Vehicle as a mobile computing platform: Opportunities and challenges," *IEEE Network*, pp. 1–1, 2023.
- [37] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [38] S. S. Srinivasa, P. Lancaster, J. Michalove, M. Schmittle, C. Summers, M. Rockett, J. R. Smith, S. Chouhury, C. Mavrogiannis, and F. Sadeghi, "MuSHR: A low-cost, open-source robotic racecar for education and research," *CoRR*, vol. abs/1908.08031, 2019.
- [39] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng *et al.*, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.
- [40] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte carlo localization: Efficient position estimation for mobile robots," *Aaai/iaai*, vol. 1999, no. 343-349, pp. 2–2, 1999.
- [41] R. C. Coulter et al., Implementation of the pure pursuit path tracking algorithm. Carnegie Mellon University, The Robotics Institute, 1992.