








Regular Article

An Autonomous Racing System: Design, Implementation, and Analysis; Team KAIST at the IAC

Chanyoung Jung^{*}, Andrea Finazzi^{*}, Hyunki Seong, Daegy Lee, Seungwook Lee, Boseong Kim, Gyuree Kang and Hyunchul Shim

Korea Advanced Institute of Science and Technology - EE Deajon Yousunggu Deahakro 291 ki c322, Daejeon 34141 Korea (the Republic of)



Figure 1. Team KAIST at the Indy Autonomous Challenge.

Abstract: While the majority of autonomous driving research has concentrated on everyday driving scenarios, further safety and performance improvements of autonomous vehicles require a focus on extreme driving conditions. In this context, autonomous racing is a new area of research that has been attracting considerable interest recently. Due to the fact that a vehicle is driven by its perception, planning, and control limits during racing, numerous research and development issues arise. This paper provides a comprehensive overview of the autonomous racing system built by team KAIST for the Indy Autonomous Challenge (IAC). Our autonomy stack consists primarily of a multi-modal perception module, a high-speed overtaking planner, a resilient control stack, and a system status manager. We present the details of all components of our autonomy solution, including algorithms, implementation, and unit test results. In addition, this paper outlines the design principles and the results of a systematical analysis. Even though our design principles are derived from the unique application domain of autonomous racing, they can also be applied to a variety of safety-critical, high-cost-of-failure robotics applications. The proposed system was integrated into a full-scale autonomous race car (Dallara AV-21) and field-tested extensively. As

^{*}These authors contributed equally to this work.

Received: 31 July 2022; revised: 5 November 2022; accepted: 8 December 2022; published: 1 June 2023.

Correspondence: Chanyoung Jung, Korea Advanced Institute of Science and Technology - EE Deajon Yousunggu Deahakro 291 ki c322, Daejeon 34141 Korea (the Republic of), Email: cy.jung@kaist.ac.kr

This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Copyright © 2023 Jung, Finazzi, Seong, Lee, Lee, Kim, Kang and Shim

DOI: <https://doi.org/10.55417/fr.2023024>

a result, team KAIST was one of three teams who qualified and participated in the official IAC race events without any accidents. Our proposed autonomous system successfully completed all missions, including overtaking at speeds of around 220 km/h in the IAC@CES2022, the world's first autonomous 1:1 head-to-head race (Figure 1).

Keywords: autonomous racing, Indy Autonomous Challenge, system design, field experiment, computer vision

1. Introduction

The Society of Automotive Engineers (SAE) defines level 5 autonomy as the full-time performance of the driving task without human intervention by an autonomous driving system. The potential advantages of autonomous vehicles are immense. According to a report by the consulting firm McKinsey and Company, autonomous vehicles will help reduce annual traffic fatalities in the United States by up to 90%, save commute time for high productivity, decrease traffic congestion and pollution, and boost the utilization of driving resources (Bertoncello & Wee, 2015). In parallel with the development of hardware technology, autonomous driving software has been a focus of intense interest among academia and industry during the past two decades. As a result, modern Advanced Driver Assistance Systems (ADAS) systems of level 3 or highly automated prototypes are readily accessible, and legislation mandating the adoption of ADAS on everyday vehicles is being pushed in certain countries (Brodsky, 2016).

To achieve further safety, reliability, and performance improvements of autonomous driving technology, the autonomous racing field of research is gaining a lot of attention these days (Betz et al., 2022). This is the same innovation pathway that links Formula 1 to our everyday vehicles. Driven by this growing interest, several real-world autonomous race events such as F1TENTH (O'Kelly, Zheng, Karthik, & Mangharam, 2020), Roborace (Rieber, Wehlan, & Allgower, 2004), Indy Autonomous Challenge (IAC, 2019), and Darpa-RACER (Darpa-RACER, 2022) have been held (see Figures 1 and 2).

Autonomous racing poses numerous research and development challenges since a vehicle is pushed to its perception, planning, and control limits (Betz et al., 2022; Talvala, Kritayakirana, & Gerdes, 2011). For example, an autonomous race car should be able to detect opponents from a far



Figure 2. Examples of real-world autonomous racing competitions. (Top left) F1 TENTH, using a 1:10 scaled vehicle. (Top right) ROBORACE, using a full-scale electric race car. (Bottom left) Indy Autonomous Challenge (IAC), using a Dallara AV-21 retrofitted Indy Lights class chassis. (Bottom right) DARPA-RACER program, using autonomous ground combat vehicles in unstructured off-road terrain at speeds.

distance and accurately predict future trajectories based on multiple factors, including environment, strategy, and race rules, in order to safely overtake. Furthermore, as professional human drivers do, autonomous racing vehicles should be able to plan their trajectory and control the vehicle at high speeds considering vehicle dynamics. Most challenging, however, is that the autonomous system should fulfill the combination of the above requirements with limited computational resources and compute everything in real time.

In this paper, we present the team KAIST's autonomy solution for the Indy Autonomous Challenge, the world's first head-to-head racing competition featuring full-scale, autonomous racing vehicles (described in Section 3). Our paper is organized into three major thematic sections: 1. Principles of the overall system design 2. A full-stack software for autonomous racing 3. Results of extensive field testing and in-depth system performance analysis. Outlining our design principles is not only helpful to understand the proposed autonomy architecture, but also to give insight into a wide range of field robotics applications. Following the design principles, we built a full-stack autonomous racing software composed of a multi-modal perception module, a high-speed overtaking planner, and a resilient control stack. Every subsystem was evaluated in the context of the autonomous racing competition. Our software stack was fully integrated into the Dallara-AV21 which served as the IAC's official vehicle platform and extensively tested in the real-world. Furthermore, we deliver the result of an in-depth computational performance analysis.

Autonomous racing has a lot of technical challenges, which come from the fact that various algorithms, including perception, planning, and control are pushed to the limits. Furthermore, system design also possesses considerable challenges. A complex autonomy stack should operate in real time within limited computing resources. Our study addressed these challenges and intensively evaluated the entire system based on real-world experiments using a full-scaled race vehicle. Contributions of this paper are not limited to the field of autonomous driving racing. We believe that our autonomy design and real-world evaluation results provide technical insight into research on various autonomous robot applications operating in extreme environments.

In summary, the contributions of this paper are the followings.

1. We propose a full-stack autonomous racing system that includes multi-modal perception, high-speed overtaking planner, and resilient control stack.
2. We integrate the proposed autonomy solution into a full-scaled autonomous race vehicle (Dallara AV-21). The result shows that our system has capabilities by reaching speeds of over 220 km/h and lateral and longitudinal accelerations of up to 6.8 and 12.4 m/s², respectively.
3. We evaluate the performance of every modules of our autonomy stack in the context of high-speed/head-to-head autonomous racing.
4. We present the system's design principles and an in-depth computational performance analysis. Even though our design principles are derived from the field of autonomous racing, we believe that these can be applied to various field robotics/safety-critical applications.

The paper is organized as follows. Section 2 presents the previous literature on the field of agile/autonomous racing. Section 3 overview of the IAC including rules, timelines, and the spec of official racing fleet. The design principles are described in Section 4. The multi-modal perception pipeline, the high-speed overtaking planner, and the resilient control stack are covered in Section 5 with a brief introduction of the official race vehicle platform, Dallara-AV21. The field testing results and the computational performance analysis are provided in Section 6. Finally, Section 7 summarizes the conclusions from this two year project and future plans for the next seasons.

2. Related Work

There are numerous existing studies relating to autonomous driving systems comprising perception, decision/planning, and control modules. This section gives an overview of previous works on the core tasks of autonomous racing systems.

2.1. Agile Control

Model Predictive Control (MPC) is a widely and well-studied controller for high-speed autonomous vehicles (Carvalho, Gao, Gray, Tseng, & Borrelli, 2013; Funke, 2015; Rosolia & Borrelli, 2019). Even though we classified MPC under the control category, a benefit of the optimization approach is that it combines some of the traditionally separate planning and control modules by creating one optimization framework that operates on a task description and dynamics model of the system. In the work of (Liniger, Domahidi, & Morari, 2015), the model predictive contouring control (MPCC), which can track a given reference path, is adapted for autonomous driving applications. The authors added a contouring cost term which represents the tracking error to the objective function. They demonstrated high-speed driving and collision avoidance experiments using a 1:43 scaled vehicle. (Goldfain et al., 2019) conducted rally racing in an outdoor, dirt environment using the Model Predictive Path Integral control (MPPI) framework, a stochastic optimal controller. The authors modeled the nonlinear vehicle dynamics using a simple neural network and trained the model in a supervised manner. They randomly rolled out the vehicle's future states repetitively and generated the optimal control signal via information theory. They integrated the proposed framework into a 1:5 scaled offroad vehicle platform and tested it in an unpaved road environment. However, the performance of model-based approaches is directly affected by the accuracy of the models. Finding the model parameters is not trivial, and the result is necessarily an approximation. Thus most previous studies simplified the optimization problem using linearized models.

With the development of machine learning technologies, recently, a number of studies have demonstrated the success of using model-free deep RL for autonomous driving, and racing (Cai, Mei, Tai, Sun, & Liu, 2020; Grigorescu, Trasnea, Cocias, & Macesanu, 2020; Jaritz, De Charette, Toromanoff, Perot, & Nashashibi, 2018; Kendall et al., 2019; Riedmiller, Montemerlo, & Dahlkamp, 2007; Wurman et al., 2022). (Fuchs, Song, Kaufmann, Scaramuzza, & Dürr, 2021) Demonstrated high-speed autonomous driving using model-free RL within a high-fidelity simulation environment by utilizing a course-progress proxy reward. They argue that because their model directly outputs the control command (one-step RL), it has the advantage of not relying on high-level trajectory planning and following while generating trajectories qualitatively similar to those chosen by the best human drivers. Even though RL approaches show promising results in the field of control, it still poses a lot of technical and practical challenges when it comes to the real-world field robotics domain. For example, RL based policy can not be clearly validated before deployment. Especially, one-step RL setup makes hard to incorporate with the other navigation and control algorithms which leads the lack of system resilience.

2.2. Planning for Racing

Given the vehicle and track models, minimum lap time or time-optimal trajectory planning has been studied in various automation fields. Multiple algorithms have been studied to create an optimal line focusing on lap-time or fuel consumption depending on their target applications. An attempt was made to develop it based on data from professional race car drivers in the racing field with the sole purpose of achieving minimum lap time (Milliken, Milliken, et al., 1995). The analysis result of the driving trajectory was used to shorten the lap time and design the vehicle. In early 2000, (Casanova, 2000) proposed a Nonlinear Programming (NLP)-based method which optimizes the path and velocity profile simultaneously using a nonlinear vehicle dynamics model. Another widely used approach is to convert the minimal lap time problem into an MPC problem (Liniger et al., 2015; Timings & Cole, 2013). They all created a nonlinear model-predictive framework to solve an optimal control problem with time as the objective. As an alternative, the shortest lap time problem can be approximated by minimizing the lateral acceleration problem, leading to finding the race line with the minimum curvature. Of course, because it does not explicitly tackle the optimization problem of lap time, this solution cannot ensure global optimality. It can, however, be useful when the precise vehicle dynamics and nonlinear

tire model parameters are unknown. (Heilmeier et al., 2019) used quadratic programming to compute a minimum-curvature racing line, and simulation testing revealed that it performs fairly similarly to the minimum-lap time. This method has the advantage of not requiring sophisticated vehicle model parameters.

Local trajectory planning for high-speed vehicles is a module that plans a noncolliding and dynamically feasible fixed horizon of trajectory near a given global race line. MPC is mainly used in the optimal control field, but it is also used as a path planner through state propagation using a model and optimal control output (Funke, Brown, Erlien, & Gerdes, 2016; Subosits & Gerdes, 2019; Williams, Drews, Goldfain, Rehg, & Theodorou, 2016). Another category for local path planning is using motion primitives in a more general control-space sampling-based planner. This approach generates multiple candidates by propagating the model given the vehicle's current state and selecting the best one based on the designed cost function. (Liniger & Lygeros, 2015) Generate a library of trajectories by forward-simulating the vehicle using a grid of vehicle velocities and steering angles up to a certain horizon. The final local trajectory was chosen by minimizing the curvature while traversing the track to maximize the velocity and keep the vehicle as straight as possible. By adjusting the size of the sampling space, this approach presents a trade-off between optimality and computational burden. To leverage the fact that the track geometry is fixed, (Stahl, Wischniewski, Betz, & Lienkamp, 2019) proposed a two steps sampling-based local path planner for a race vehicle. In an offline process, they build the track as a graph consisting of layers, nodes, and edges. The connectivity of node and edge is determined based on the kinematic constraints of the vehicle. Then, a cost is assigned to each edge according to the displacement from the race line and the curvature. The local path is planned in the online process by performing a minimum cost path search of a fixed time horizon branch based on the track graph model. This planner was integrated into a Roborace vehicle and validated over the 200 km/h speed range.

2.3. Autonomous Race Vehicles and Competitions

The majority of research on autonomous racing has been conducted utilizing simulation, or scaled vehicle platforms (Herman et al., 2021; Kabzan et al., 2020; Liniger & Lygeros, 2017; Weiss & Behl, 2020). Exploiting advanced physics engine and graphic technology, simulation enables a variety of experiments that would be difficult to carry out in the real world. Also, considering the cost of operation and safety, the scaled platform has been employed in several researches. Unfortunately, there is still a gap between simulation and reality, making simulation impossible to estimate the algorithm's scalability precisely. This section will concentrate on a full-scaled autonomous vehicle and competitions.

In collaboration with Stanford University, Audi debuted Shelly, an autonomous TTS capable of high-speed autonomous driving (Funke et al., 2012). They reached a top speed of about 305 km/h at the Salt Flats in Utah, and high-speed autonomous driving at the Pikes Peak International Hill Climb in 2009. Since they focused on time trial race or high-speed driving, it is only equipped with an integrated Differential Global Positioning System (DGPS) and Inertial Measurement Unit (IMU) as sensors.

The Roborace platform is based on an electric vehicle designed for the autonomous racing competition (Rieber et al., 2004). The platform is based on a Le Mans Prototype chassis and is equipped with cameras, LiDARs, and radars. Starting in 2018, the competition was held as a season event, with time trials. There were virtual obstacles on the track, and each team should develop a software stack aiming to complete the race mission without any time penalty. Vehicles are not provided for each team but rather in a way that allows their software to be deployed across vehicles.

The Indy Autonomous Challenge (IAC) is the most recent autonomous racing competition. The Dallara-AV21, the official vehicle platform for the competition, was constructed on an Indy Lights chassis and has a combustion engine. More information about the race will be provided in the section that follows.



Figure 3. Overview of the IAC timeline. (Left) Simulation phase. (Middle) The first real-world high-speed competition at IMS. (Right) 1:1 head-to-head autonomous race event at LVMS.

3. Indy Autonomous Challenge

As shown in Figure 3, IAC began in November 2019. It was divided into several stages: document screening, hackathons, simulation races, and real-world races. About 30 teams participated in the simulation race, and 9 of them, including our team, competed in the real-world autonomous race competition. The first real-world competition held at Indianapolis Motor Speedway (IMS) in 2021 and was conducted with a time trial racing format. Teams had to complete pit-in, pit-out, performance lap, and static avoidance tasks following the race control signals. The next race, part of the official CES 2022 event, was held in a 1:1 head-to-head race format, with two vehicles racing on the course simultaneously. Several artificial rules were imposed: 1. The defending vehicle must only run on the inner line of the track at the commanded speed. 2. The defending and attacking vehicles must adhere to race control, and overtaking is restricted to a specific region and race flag. 3. The attacking vehicle should close the overtaking maneuver when a safety distance of 20 m is secured.

The participating teams used the same vehicle platform, the Dallara AV-21, adapted for autonomous racing, and they had to design and integrate their own autonomy stack. The vehicle is rear-wheel drive, powered by an internal combustion engine which produces 335 kW (449 hp) and has a six-speed sequential gearbox. Computing devices, sensors, and controllers were placed instead of the driver seat. Six Gig-E cameras, three Radars, three solid-state LiDARs, and an RTK GPS are equipped as a sensor package. The computing platform included an Intel Xeon CPU with an Nvidia Quadro RTX 8000 GPU. Figure 4 shows the system diagram of the Dallara AV-21.

4. Design Principles

Autonomous racing poses a few challenges in addition to the functional requirements of an urban-oriented autonomous driving system. The capabilities of the platform in terms of acceleration and maximum speed, along with the particular environment that are racing tracks, make the application domain unique. Moreover, the pioneering nature of the IAC competition introduces a high level of uncertainty in the definition of the functional requirements.

To cope with the above-mentioned challenges, dependability, evolvability, and performance were identified as the highest-priority software qualities at the early stages of the design process. Our key design principles are shown in Figure 5. Dependability is the measure of the trustworthiness of a software system. (Sommerville, 2015) describes it as a five-dimensional quality. Being autonomous racing a safety-critical and high cost-of-failure application, we define dependability as an aggregate measure of reliability, safety, and resilience, as availability and security are of secondary importance in this work's application domain. Evolvability refers to the ability of a system to respond to functional requirements modifications, handle domain uncertainty, and absorb change without incurring disruption. This software quality is an important aspect of field-robotics systems since the requirements of the problem can change according to various factors, such as environmental

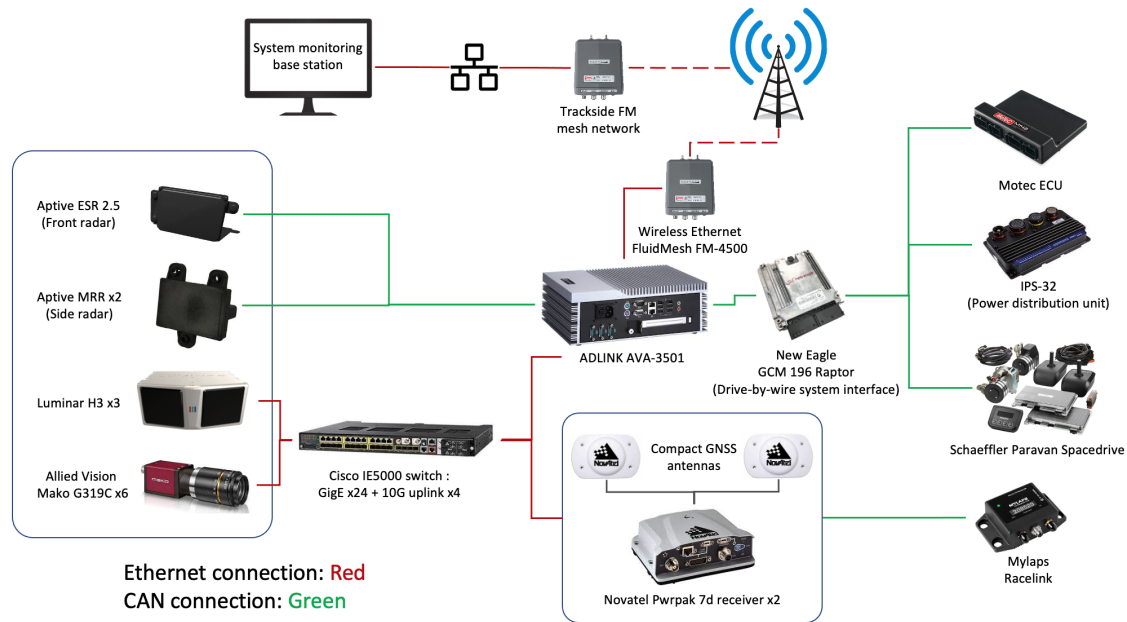


Figure 4. System diagram of the Dallara AV-21.

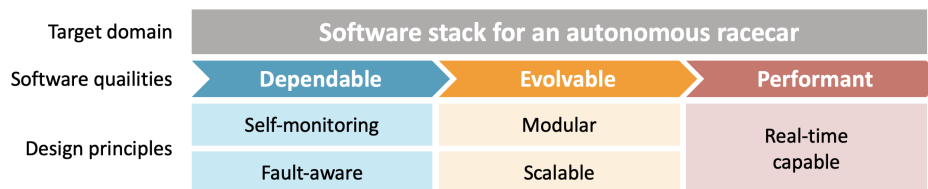


Figure 5. Key design principles of team KAIST autonomy stack.

conditions, field testing scenarios, and so on. Lastly, performance; this can be defined in several ways, and different classifications of performance and real-time systems are available in the literature (Oshana, 2006). In our work, performance is intended as the capability of the system to respect execution deadlines and operation frequencies defined at the design stage, analogously to what is asked to a soft real-time system.

Out of these properties, the following requirements were defined and adopted throughout the development of the system.

- **Self-monitoring:** the system has the ability to detect and react to potential failures at different levels (infrastructure, application, platform).
- **Fault-aware:** the potential occurrence of faults (either internal or external) is taken into account as a fundamental modeling principle. Clear separation of duties and a fail-fast (Shore, 2004) development approach were adopted to satisfy this requirement.
- **Modular:** adding, removing, and exchanging functional components with minimal integration effort is necessary to respond effectively to software evolution and functional requirements uncertainty.
- **Scalable:** the architecture as a whole, as well as the single components, must be able to grow as the system's complexity increases and new functionalities are introduced.
- **Real-time capable:** execution deadlines and operation frequency have to be defined and respected by the software components, as in an embedded soft real-time system.

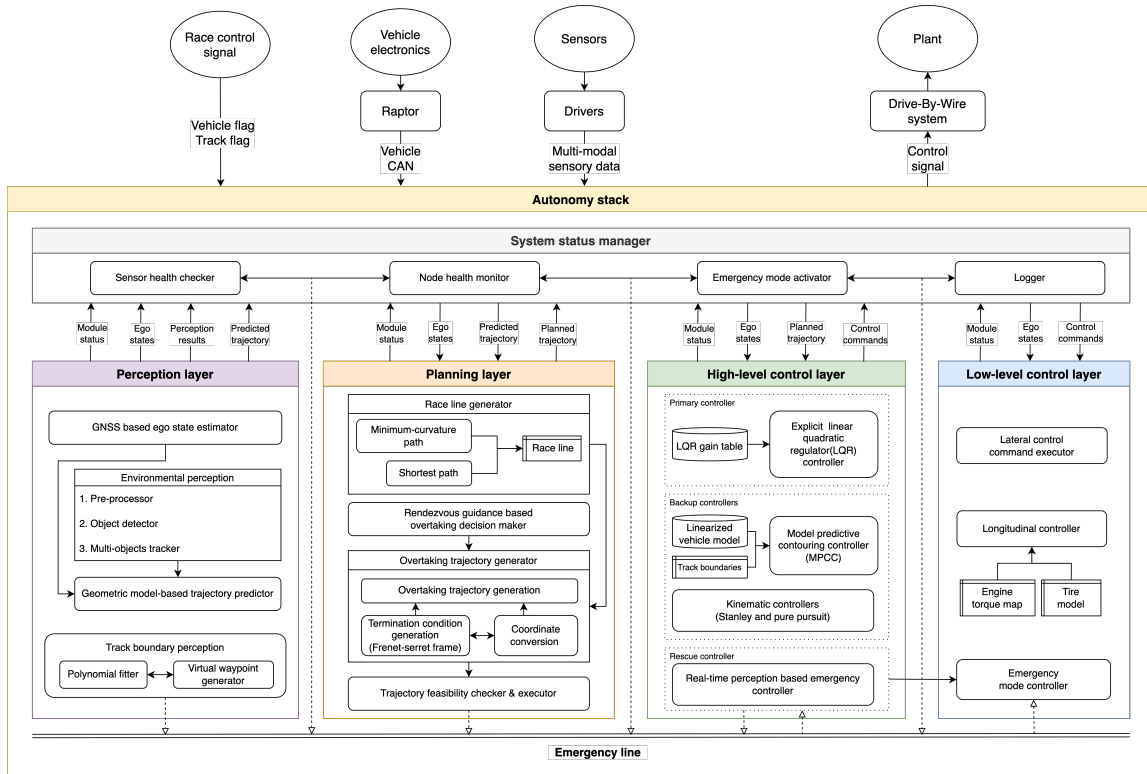


Figure 6. Overview of team KAIST’s software architecture.

5. Autonomous Racing System

5.1. Architecture Overview

Figure 6 depicts a high-level functional block diagram of the proposed autonomous racing system. Our autonomy stack is designed to address the head-to-head race scenarios. The system comprises several subsystems: (1) system status manager, (2) perception layer, (3) planning layer, and (4) high-level and low-level control layers. Every subsystem is developed following our design principles with different requirements (e.g., update frequency, type of output).

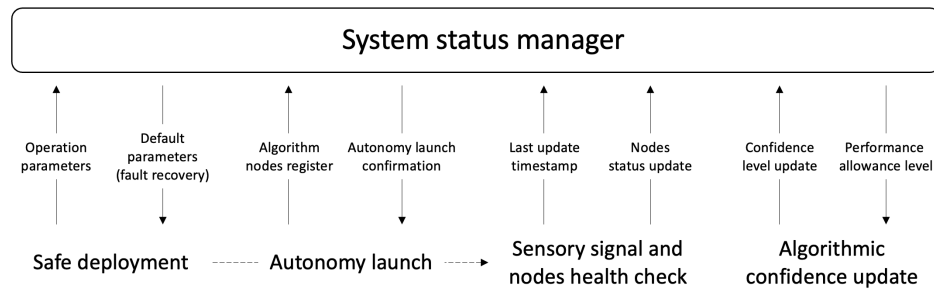
As previously stated, our autonomy stack follows the traditional perception-planning-control task flow. We designed a multi-modal perception pipeline based on the equipped sensors of the Dallara-AV21. The perception module’s outputs are fed into the planning module. Then, the planning module is responsible for generating a collision-free trajectory that can also overtake the other opponent. Finally, the control subsystem computes the desired lateral and longitudinal control commands to precisely follow the planned trajectory while considering vehicle dynamics. Furthermore, our control stack is designed to support multiple control algorithms running in parallel to cope with system failures. Besides the functional subsystems (perception, planning, and control layers), a system status manager was designed to ensure system resilience by monitoring the health of each subsystem. For instance, if any module is not in the nominal status, the system status manager dynamically reacts at a system level to recover or safely stop the vehicle. More information regarding our system will be provided in the following sections.

5.2. System Status Manager

Autonomous racing is a safety-critical and high-cost of failure application. As a result, a fully automated mechanism is necessary to detect system abnormalities and initiate appropriate recovery

Table 1. Pre-set node status list and codes.

Node Status Name	Code
NODE_OK	0
NODE_INITIALIZED	64
NODE_INACTIVE	100
NODE_NOT_INITIALIZED	128
NODE_ERROR	200
NODE_DEAD	255

**Figure 7.** Overview of the behavior of system status monitor by operation stage.

behavior. With that aim, we have configured a system status manager (hereafter referred to as SSM) that supervises the state of submodules that compose the autonomy stack. As shown in Figure 7, our SSM contributes to the system’s resilience through different criteria according to the operation phase.

The SSM starts by checking parameters in the operation configuration to ensure safe deployment. Multiple parameters must be preset according to the testing purpose and race operation strategy. However, manually inputted parameters by the user may be incorrect and lead to significant performance degradation or fatal accidents during vehicle operation. To prevent faulty operation by human operators, our SSM validates the parameters during the deployment/launching phase, such as preset value range, sign, and data type. Representative parameters managed by the SSM include maximum speed, maximum acceleration/deceleration, time-out threshold for sensors, and watchdog threshold. In case any parameter does not meet the conditions, the SSM immediately suspends the autonomy launch and applies full braking.

After the SSM confirms that configurations are OK, every algorithm module is launched. During the online autonomy phase, the SSM periodically checks the health of sensory signals. Here, we consider the update rate and the size of the sensor data as health indicators. The SSM manages all the status of sensory data centrally, and every algorithm module which utilizes the raw sensor data was strictly designed to check its health before the algorithmic computation. Similarly, all algorithm blocks (hereafter referred to as nodes) that compose the autonomy stack report the node status to the SSM. As shown in Table 1, the state of each node is updated based on its phase. As a result, the SSM determines that the current system is under a nominal situation only when all registered node status codes are OK (code 0). The SSM responds differently depending on the type of error and on the node where the error occurs.

To monitor the system’s state more precisely, the nodes communicate the reliability of algorithm’s result during operation. Measuring the reliability or uncertainty of solutions is one of the key aspects of designing a resilient system. Our idea was to cope with abnormality by using the degraded mode of operation. For example, when our system is not confident about the GNSS-based localization result, our SSM lowers the maximum speed, acceleration, and so on. Furthermore, when the localization node incurs into fatal error at high-speed driving, the SSM switches to a controller that does not rely on GNSS-based location signals. Details about our efforts for monitoring the system status will be provided in the according sections.

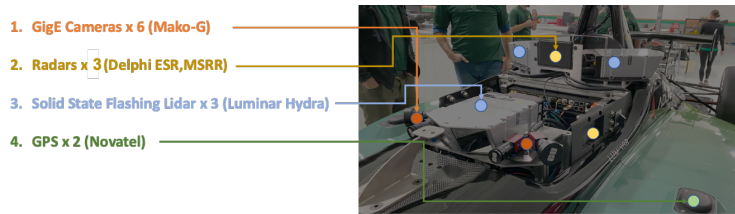


Figure 8. Dallara AV-21 sensor configuration.

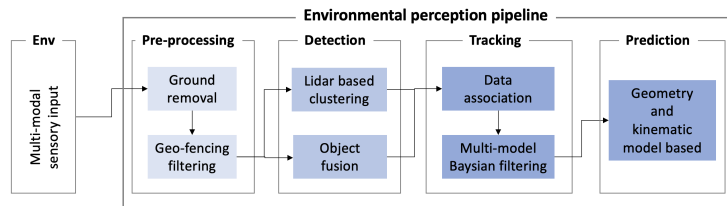


Figure 9. Schematic flow diagram for the object detection pipeline.

5.3. Multi-modal Perception Pipeline

Estimating the ego’s state and understanding the surroundings are the first steps of the autonomous driving task. Perception performance directly impacts the system’s capability and its overall robustness. Unique features of the perception system for autonomous racing are as follows. (1) Since the race vehicle drives only in a known and controlled environment (i.e., a race track), the perception system can fully utilize its geometric information. (2) It can be reasonably assumed that only race vehicles exist on the track in most cases. (3) Considering the driving speed, the perception range and the detection update rate should be long and fast enough to safely operate.

Given the nature of high-speed autonomous driving problem, we built a multi-modal perception pipeline using the sensory system of the vehicle, as shown in Figure 8. Our perception stack mainly comprises a GNSS-based ego state estimator and an environmental perception part, which includes detection, tracking, and prediction (see Figure 9). In the following sections, we will present the details of our approach and our implementation results.

5.3.1. Ego State Estimation

Reliable state estimation is crucial for autonomous robots, especially under high-speed driving conditions. Since it is a long-studied problem, there are various ways, including Simultaneous Localization And Mapping (SLAM), odometry estimation, and GNSS-based navigation. Each approach presents advantages and weaknesses in accuracy, consistency, computational burden, etc. However, after testing SLAM-based methods (Qin et al., 2020; Shan & Englot, 2018; Shan et al., 2020) at the IMS, we concluded that SLAM-based navigation is not appropriate for the racing domain since it needs considerable computation resources. Also, it does not work well in feature-poor environments such as long straight sectors.

Eventually, we designed our state estimator using two GNSS receivers with Kalman Filtering (KF). Our state estimator was designed to output the state information and an indicator of the solution quality for system-level resilience. The motivation behind this design choice is that, if one of the two GPS units has critical issues due to failures in signal receiving (e.g., spikes, drift, null) or hardware malfunctioning, the state-estimation task can be carried out relying on the other unit (Failure of both GPS units case should be handled differently and we will introduce our solution in Section 5.5.3).

Figure 10a depicts our state estimator architecture. Two KF-based state estimation algorithms (Karimipour & Dinavahi, 2015) are running in parallel using two different GNSS sources. Each

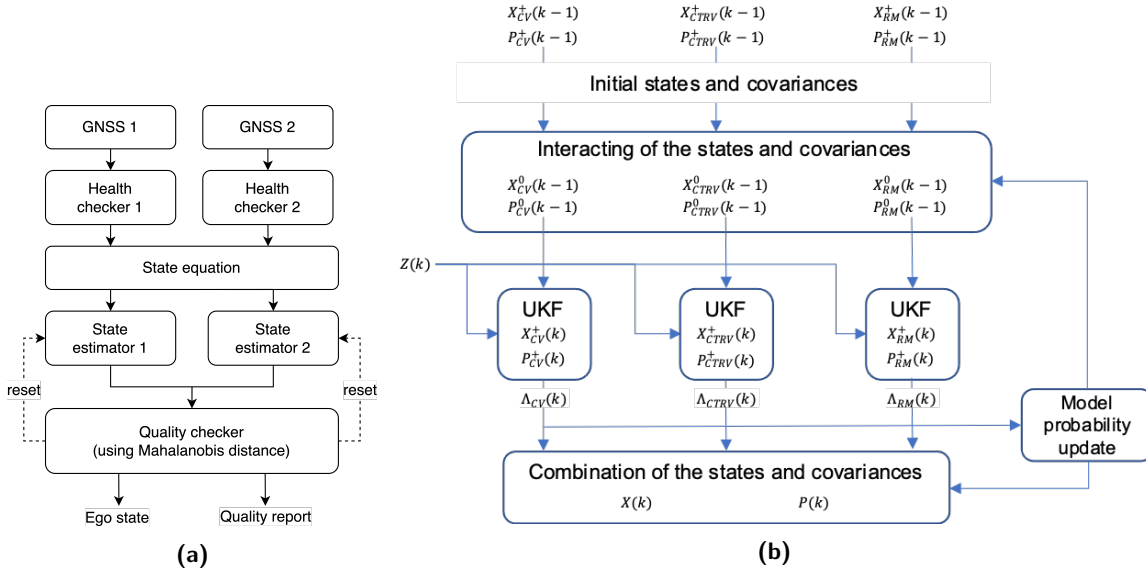


Figure 10. (a) Resilient ego state estimation pipeline. (b) Schematics of one cycle of the IMM-UKF-PDAF using CV, CTRV, and RM models.

algorithm estimates position, orientation, velocity, and acceleration. Two estimation outputs are passed to the quality checker. The quality checking part has two main functions related to reliable state estimation and system-level resilience: (1) measures the quality for the individual estimation outputs and selects the best and (2) reports the state estimator’s health to the system status manager. We measured the quality of estimation using a Mahalanobis distance (De Maesschalck, Jouan-Rimbaud, & Massart, 2000) between estimation and sensor measurements as follows:

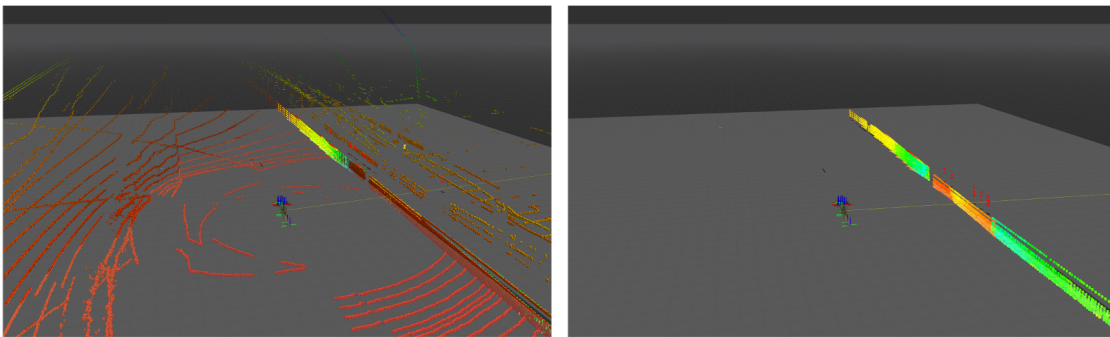
$$D_h \triangleq (\mathbf{x} - \mathbf{z})^T \Sigma^{-1} (\mathbf{x} - \mathbf{z}), \quad (1)$$

where \mathbf{x} is the estimation output, \mathbf{z} is the sensor measurement, and Σ^{-1} is the process covariance. We empirically set some ranges of Mahalanobis distance to represent the state estimation quality and is reported to the system status manager. For more details of our localization module, please refer to (Lee, Jung, Finazzi, Seong, & Shim, 2022).

5.3.2. Opponent Detection and Tracking

LiDARs and radars were used for opponent vehicle detection. Even though we had implemented a vision-based deep learning detection algorithms (Howard et al., 2019; Liu et al., 2016), a few practical issues prevented us from using this method. First, we could not collect enough vision data in the real-world environment for training and validation. As an alternative, we augmented our training dataset using web-crawled data and a simulator. However, the real-world detection performance was not reliable enough. Also, external sensor fusion between cameras and LiDARs was not fast enough for high-speed driving scenarios.

Our perception pipeline begins with preprocessing sensor data. The ground removal algorithm distinguishes between the ground plane and nonground points in lidar point cloud data. The most representative process is to assume one plane within the sensing range and remove the points associated with the plane using plane fitting. However, the target track is an oval with banking in all sectors, and the change in bank angle is significant, particularly in the area entering or exiting the corners from or into the stretch. To that end, we created a nonground point filtering algorithm based on height distribution (see Algorithm 1). Our preprocessing method projects three-dimensional LiDAR points into XY plane and represents into grid cells. After that, two factors are considered: (1) number of points in the grid and (2) Height distributions of the grid. When these two different

Algorithm 1. Height-distribution-based ground filtering algorithm.**Require:** Raw LiDAR point cloud, P , Grid cell size, g_c **Ensure:** Non-ground LiDAR points, P_n $P_p \leftarrow \text{projectToXYplane}(P)$ $G \leftarrow \text{gridDiscretization}(P_p, g_c)$ $P_n \leftarrow \{\}$ **for** $g \in G$ **do** $n_p \leftarrow \text{countNumOfPoint}(g)$ $h_d \leftarrow \text{calcHeightDistribution}(G, g)$ **if** $n_p > n_{thres}$ and $h_d > h_{thres}$ **then** $R_n \leftarrow \text{reconstructTo3D}(G, g)$ $P_n = P_n \cup R_n$ **end if****end for****return** P_n **Figure 11.** Ground filtering result of 3D point clouds based on height distribution. (Left) Merged LiDARs point clouds before ground filtering. (Right) Ground filtered result.

criteria are over thresholds, we assumed that the grid is occupied by nonground objects. Figure 11 shows the output of ground filtering on the LiDAR point cloud. In addition, a geofencing filter is applied using a track model for both LiDAR and radar data. Sensor data around 1 m from the track boundary was removed.

After the preprocessing step, the filtered data contains only a few points corresponding to the opponent vehicle on the track. Then, we adapted a clustering algorithm that is widely used in autonomous robots for grouping and labeling points associated with an object (Uppada, 2014). The clustering of point clouds in this work is accomplished using a hierarchy-based method (Madhulatha, 2012). The hierarchy-based clustering algorithm, also known as connectivity-based clustering, classifies objects based on the distance between neighboring points. However, this clustering method is sensitive to outlier points and may result in separating a single object into multiple clusters. We used two consecutive clustering steps with different connectivity criteria to obtain more accurate results. The clustering result was defined as the object's center point and maximum width and length. By comparing the finally detected width and length with the vehicle's geometry, we cross-checked the opponent's relative position on the track.

Clustered objects represented in 3-Dim local coordinate are fed into the tracking module to estimate the object's dynamic states. This information is necessary for the trajectory prediction module, the last step of our environmental perception module. The tracking problem can be modeled as a filtering problem in which the object states might be noisy. Bayes filtering is one of the widely used statistical theories that can be applied directly with multiple types of models. Because of its heuristic-free approach, this filter can generally be applied to numerous robot applications (Thrun, 2002). Kalman Filter (Gutman & Velger, 1990) is the analytical implementation of the Bayesian

method that seeks to compute the optimal filter gain from its posterior density recursively. Generally, it assumes that the target objects' dynamic and posterior density at the previous observation follows Gaussian distribution, and the measurement function is linear. The Extended Kalman Filter (EKF), as well as the Unscented Kalman Filter (UKF) (Wan, Van Der Merwe, & Haykin, 2001), are filters developed to overcome the limitation imposed by linear modeling when trying to capture object motion.

We implemented the Interacting Multiple Model UKF Probabilistic Data Association Filter (IMM-UKF-PDAF, or IMM for short) proposed by (Arya Senna Abdul Rachman, 2017) in order to track robustly in varied race circumstances where $X(k)$, $P(k)$, $Z(k)$, and $\Lambda(k)$ represent state, covariance, measurement, and the likelihood for the observation at time k , respectively. It estimates an object's ambiguous dynamic behavior by combining several models rather than using a single motion model for the existing filter for state estimation. IMM can be made up of n filters that run different models in parallel and output individual probabilities. Then, IMM uses a weighted average of each model output to calculate a single combined estimate state and its corresponding variance for the next iteration. We employ three models for object tracking: (1) constant velocity (CV) and (2) constant turn rate and velocity (CTRV) (3) Random Motion (RM). We set the initial weights for each model to 0.5, 0.3, and 0.2, respectively. Figure 10b schematically shows one cycle of our tracking module's operation.

Finally, the output of the tracking module contains the 3-Dim position and XY velocity vector in local coordinates. The results are transformed into global coordinates using the ego vehicle's status and upsampled to 50 Hz. Figure 12 illustrates the detection performance evaluation result during the semifinal race with *PoliMOVE* (For evaluation, we used 50 samples from the logged data during the race event at LVMS. *PoliMOVE* provided their GPS log, which we used as ground truth.). The LiDAR-based perception result showed the highest average precision (AP); however, its maximum detection range was around 70 m. On the other hand, raw radar data showed the longest detection range but lowest recall performance. Our multi-modal detection approach, as shown in the green, showed a balanced performance in terms of accuracy and range. Additionally, we conducted the perception performance analysis according to the range. We used the L2-norm distance to measure how our detection output is geometrically close to the ground truth. Note that, we only collected samples when the driving speed is over 100 mph. For generating the ground truth, we manually measured the position of the other vehicle using ego-vehicle's state and raw Lidar data. We divided the validation samples into three cases: closer than 50 m, over 50 m but less than 100 m, over 100 m but less than 120 m. As shown in the Figure 12 right, our deployed method (Lidar and radar based detection and tracking) showed a well balanced in terms of detection accuracy and its performance consistency. Given the validation case when the sampled targets are closer

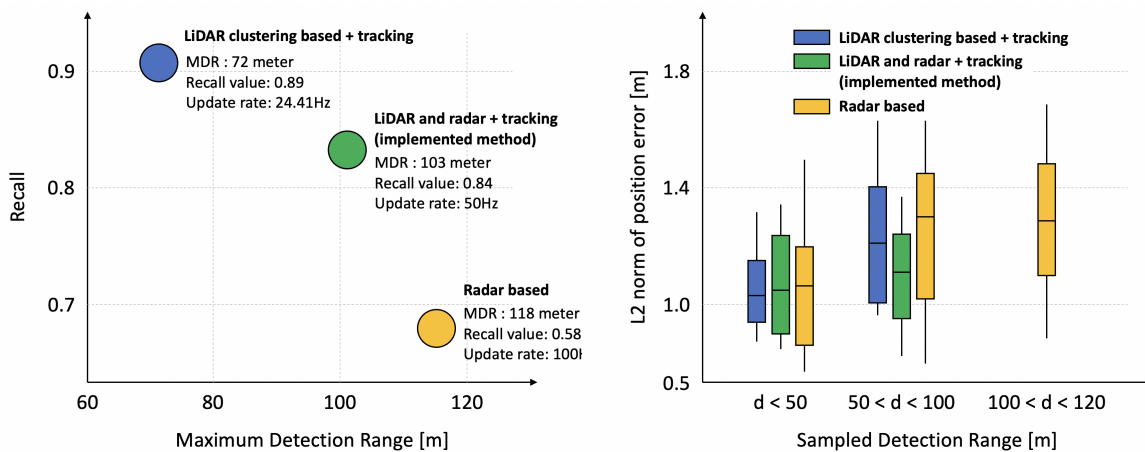


Figure 12. Perception performance analysis results. (Left) AP and inference speed comparison results for three different methods. (Right) L2 norm of position error according to sampled detection ranges.

than 50 m, Lidar only clustering based detection method showed the best performance among three implementations.

5.3.3. Future Trajectory Prediction

Based on the perception results, autonomous vehicles must plan a safe motion trajectory considering the predicted opponent's trajectory. Accurate prediction of other vehicles' trajectories increases the safety of the autonomous vehicle and has an impact on general traffic safety, and efficiency (Claussmann, Revilloud, Gruyer, & Glaser, 2019). For these reasons, substantial research on trajectory prediction and ego-motion planning has been actively conducted in autonomous mobile robotics (Elbanhawi & Simic, 2014; Laumond et al., 1998; Petti & Fraichard, 2005).

The majority of predictions for surrounding vehicles were made in an urban environment. However, in a racing scenario, the opponents' trajectory prediction differs from that of an urban setting. In an urban driving environment, for example, road geometry is one of the essential cues for the prediction module. In a race, however, the race line, which shows the minimum lap time, can be used for predicting an opponent's future trajectory. Furthermore, a strategic trajectory that prevents neighboring vehicles from overtaking can predict a free-racing scenario (one in which no arbitrary race restrictions apply). A few studies have modeled race as a noncooperative game. In our previous work (Jung, Lee, Seong, Finazzi, & Shim, 2021), we used the Stackelberg Game to model free-racing. We consecutively build independent Stackelberg games amongst neighboring vehicles to change the n-player game to two-player games, and each game is solved recursively. We assumed in each game that two cars were seeking to maximize their progress, with the following agent accounting for collisions. The progress term is converted to a payoff function similar to the MPC objective function. Our method was evaluated in a simulated environment and demonstrated the capability to pass in multi-agent competitive race scenarios.

However, the real-world IAC competition was ultimately decided to be a 1:1 overtaking competition for safety reasons. Also, the race rules imposed the defending vehicle to maintain a commanded speed (by racing control) along the inner side of the track. We designed the prediction task more straightforwardly following the competition's rules, and we made the following two reasonable assumptions for prediction:

1. The opponent will stay on track $\mathbf{X}_{\text{track}}$.
2. During the prediction, the opponent will keep lateral displacement from the track center line and velocity.

The prediction problem is to forecast the set of future states $P_{i=1,2,\dots,n} = \{p_i^0, p_i^1, \dots, p_i^{t_{\text{pred}}}\}$ for the prediction time instant 0 to t_{pred} , where $p_i^t = (x_i^t, y_i^t, v_i^t)$. x, y , and v denote the spatial Cartesian coordinates x and y and velocity, respectively. We used the Frenet coordinate system for prediction which is a well-known coordinate system in trajectory planning and control theory. The Frenet frame is made up of three vectors: the normal vector \vec{n}_r , the tangential vector \vec{t}_r , and the binormal vector \vec{b}_r . These vectors can describe the kinematic properties of a particle moving along a continuous curve, as illustrated in Figure 13, $s(t)$ and $d(t)$ represents the progress alongside the reference path and the lateral displacement at time t , respectively.

Assuming that the track's center line is \vec{r}_c , the predicted trajectory of the surrounding vehicles \vec{x} is expressed in Frenet coordinate space as follows, following the aforementioned assumptions (see Figure 13):

$$\vec{x}(s(t), d(t)) = \vec{r}_c(s(t)) + d(t)\vec{n}_c(s(t)); \quad (2)$$

1. $\vec{x}(s(t), d(t)) \in \mathbf{X}_{\text{track}}$;
2. $d(t) = d(t_0)$, $\dot{s}(t) = \dot{s}(t_0)$, $\ddot{s}(t) = 0 \forall 0 \leq t < t_{\text{pred}}$.

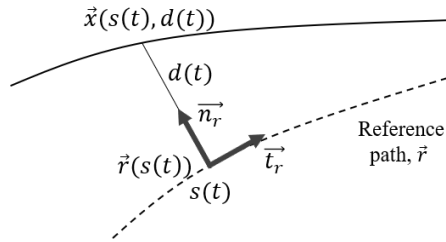
We calculated the lateral displacement from the center line of the track, $d_0 = d(t_0)$, and its progress, $s_0 = s(t_0)$. Here, we set the prediction time horizon to 4 seconds with a sampling time of

Algorithm 2. Prediction solution updating algorithm.

```

1: procedure updatePrediction
2:   while True do
3:      $t_p \leftarrow \text{GetCurrentTime}()$  ▷ Update time.
4:      $F_p \leftarrow \text{CheckPerceptionOutput}()$  ▷ Check whether there is perception output.
5:     if  $F_p$  then
6:        $P_c \leftarrow \text{GetPerceptionOutput}()$  ▷ This only activate when there is perception result.
7:        $t_c \leftarrow \text{GetCurrentTime}()$  ▷ Update time when the perception result is updated.
8:        $P_p \leftarrow \text{GenPredictionOutput}(P_c)$ 
9:       return( $P_p$ )
10:    else
11:      if  $P_p \neq \emptyset$  and  $t_c + t_{thres} \leq t_p$  then
12:         $P_p \leftarrow \text{PropagateOnPreviousResult}(P_p)$ 
13:        return( $P_p$ )
14:      else
15:         $P_p \leftarrow \text{Clear}()$ 
16:      end if
17:    end if
18:  end while
19: end procedure

```

**Figure 13.** Representation of a trajectory in a Frenet frame.

0.1 sec. Prediction results are kept even if perception results are not updated for a certain period (here, we set 2 secs) to cope with the instant loss of the tracked objects. To this end, the internal loop updates the trajectory by propagating the object's status using the previous prediction result. Algorithm 2 shows the principal steps of our prediction process.

5.4. High-speed Overtaking Planner

5.4.1. Global Race Line Generation

The definition of the race line is the path that minimizes the lap time considering the vehicle dynamics and track geometry. Following this definition, the race line generation problem can be modeled as a nonlinear optimization problem as follows:

$$\begin{aligned}
 & \text{Minimize } t = \int dt = \int \frac{dt}{ds} ds = \int \frac{1}{v} ds \\
 & \text{subject to } \kappa v^2 - \mu g \leq 0, \\
 & \quad \kappa \leq \kappa_{\max}, \\
 & \quad v \leq v_{\max}, \\
 & \quad a_{\min} \leq a \leq a_{\max},
 \end{aligned} \tag{3}$$

where t is the total time, s is the vehicle's travel length, v is the vehicle's velocity, a is the acceleration, μ is the tire-road friction coefficient, and κ denotes the curvature.

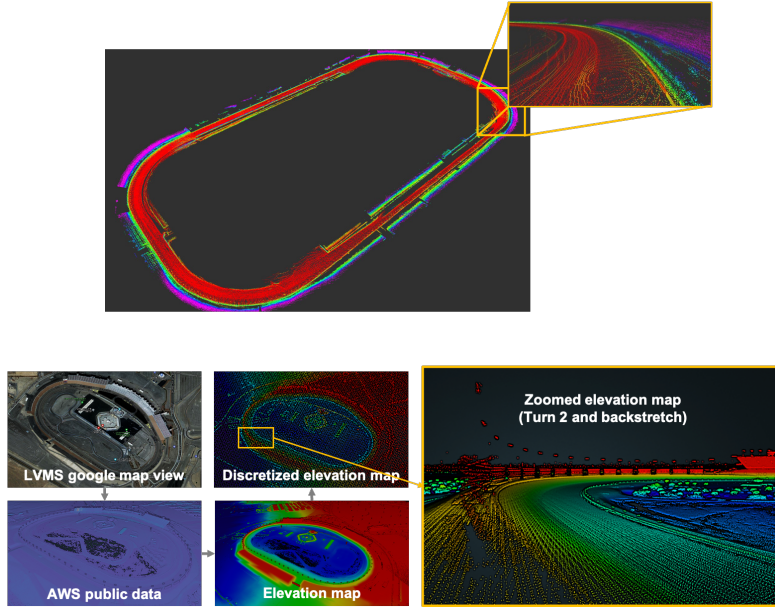


Figure 14. Visualization of track modeling results. (Top row) Indianapolis Motor Speedway (IMS) mapping result. (Bottom row) Las Vegas Motor Speedway(LVMS) track modeling flow using USGS data.

Various approaches to calculating the optimal race line have been proposed (Christ, Wischnewski, Heilmeier, & Lohmann, 2021; Dal Bianco, Bertolazzi, Biral, & Massaro, 2019; Lovato & Massaro, 2022). Most of the previous researches utilized nonlinear vehicle and tire dynamic models and solved Equation 3 directly via optimization framework, and the properties of the vehicle were incorporated into the constraint functions. However, acquiring the model parameters is not trivial. Also, the performance degradation due to model mismatch is one of the well-known limitations of “model-based” approaches. (Heilmeier et al., 2019) demonstrated experimentally that the minimum curvature trajectory performs nearly as well as the model-based optimal racing line. The path’s curvature is minimized, which maximizes the achievable velocity.

We generated the race line using the minimum curvature trajectory following the idea from (Heilmeier et al., 2019). On top of that, we also considered the travel distance since the lap time is a function of speed and distance. To this end, we started from collecting the track boundaries. We modeled the track using LEGO LOAM (Shan & Englot, 2018) and the United States Geological Survey (USGS) LiDAR point cloud data. We manually extracted the track’s inner and outer boundary points from the 3D map and interpolated them using a cubic spline model. Figure 14 shows our map-building results and examples of USGS data. After that, we generated the geometrically centered line of the track, the reference line representing the track, using the average operation on the sampled two-track boundary point sets, and it can be written as follows:

$$\vec{r}_n = \vec{p}_n + \alpha_n \vec{n}_n, \quad (4)$$

where \vec{p}_n is the center line point, \vec{n}_n is the unit normal vector, and α_n is the distance to the track boundary.

Given the track model, minimum curvature trajectory was generated using open source implementation from (Heilmeier et al., 2019). Furthermore, the shortest path was incorporated using a simple geometric-weighted sum operation. Assuming that the minimum curvature race line is P_{mc} and the shortest travel distance path is P_{sd} , our final race line can be represented as follows:

$$P_c = \zeta P_{mc} + (1 - \zeta) P_{sd}, \quad (5)$$

Table 2. Results of estimated lap times under different weights, ζ , and maximum speed limit conditions. If the maximum speed of the autonomous system is below the dynamics limit of the vehicle itself, it can be seen that the minimum curve-based race line does not always show the minimum lap time.

Track	Maximum Velocity Constraint	Estimated Lap Time [sec]			
		$\zeta = 0.0$	$\zeta = 0.3$	$\zeta = 0.6$	$\zeta = 0.8$
LVMS	220 km/h	44.003	43.197	42.847	42.889
	170 km/h	56.266	56.149	56.423	56.501

where ζ is the geometric weight parameter between two different paths. Note that, since our target track is an oval track, we set the inner boundary as the shortest path.

We verified our idea by experimentally measuring the lap time of P_c generated under various conditions using a high-fidelity racing simulator. Table 2 shows the estimated lap time results. Note that estimated lap times can be different from the real-world testing results. However, we can see that incorporating the shortest path into the race line can help to reduce the lap time in certain operation conditions (e.g., maximum speed).

5.4.2. Local Trajectory Planning for High-speed Overtaking

Overtaking is referred to as the core of racing. Professional race car drivers simultaneously push the vehicle to the limit while adhering to the race rules and performing strategic moves based on contextual understandings of opponents' intentions. The local trajectory planner takes responsibility for high-speed overtaking in our autonomy stack.

The proposed local trajectory planner is primarily made up of three hierarchical module stages: (1) decision, (2) planning, and (3) feasibility checking. The predicted trajectory of the opponent, the current ego states and race flag (provided from the race control), and the planned trajectory from the previous step were used as inputs. Our planner starts by making overtaking decisions based on the current ego state and the predicted opponent's trajectory. The decision to attempt overtaking or not is crucial in high-speed driving scenarios since it can not be accomplished even though it is collision-free. To decide whether overtaking can be safely done under our operation setup and system capability, we adapted the Rendezvous Guidance (RG) (Kunwar & Benhabib, 2006) and applied it to the overtaking problem. The decision module determines whether overtaking is possible within a given time frame. Then, the decision outcome is passed to the sampling-based trajectory generation module. Our trajectory generation module plans the optimal jerkiness trajectory (Werling, Ziegler, Kammel, & Thrun, 2010) considering the driving speed and the vehicle dynamics. As the last step, the collision and tire model-based feasibility were checked before execution. Figure 15 depicts the overall structure of our hierarchical planning module. In the following, we will deliver details of each step.

Rendezvous Guidance based Overtaking Decision Making. The proposed overtaking decision-making algorithm is inspired by (Ghumman, Kunwar, Benhabib, et al., 2008) and is based on the RG (Kunwar & Benhabib, 2006). The rendezvous problem is well-known as a chase-target problem in which the target does not require evading acceleration due to common space debris. RG was used for the space docking mission. Analytically, it has been demonstrated that RG provides the optimal solution for rendezvous with immobile targets based on the parallel navigation law. There are two types of participants in an autonomous rendezvous: chaser and target (see Figure 16a). Given current chaser's velocity, v_c , and target vehicle's velocity, v_t , RG is responsible for generating the velocity control signal of the chaser that will eventually match these two plants. According to the parallel-navigation law, the relative velocity between the chaser and the target, \vec{r}' , should remain parallel to the Line of Sight (LOS), \vec{r} . If this condition is satisfied, the chaser's distance from the target will decrease until they collide.

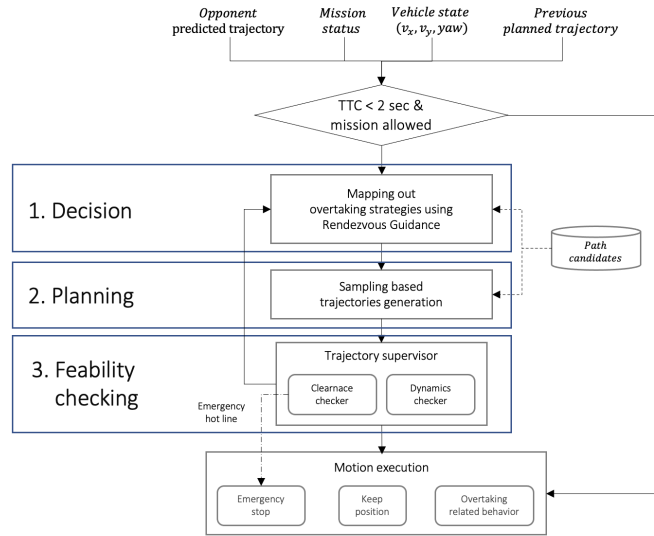


Figure 15. Schematic diagram of hierarchical overtaking trajectory planning module.

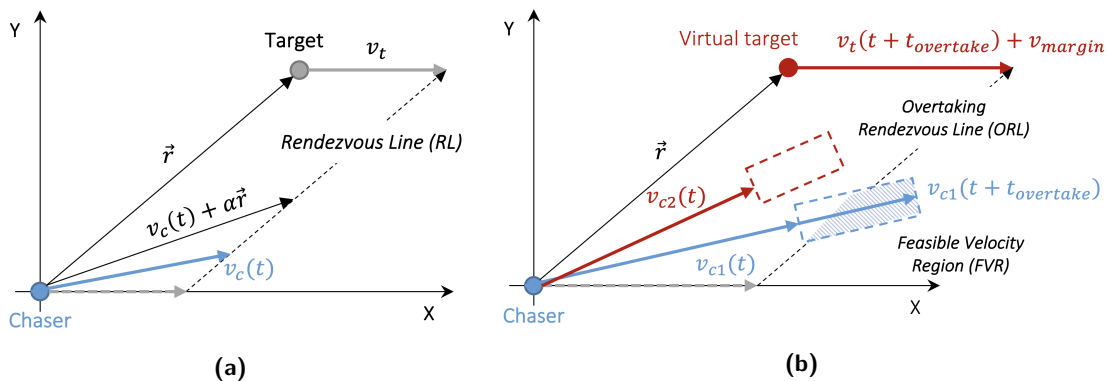


Figure 16. (a) Construction of rendezvous line based on parallel-navigation law. (b) Construction of ORL and FVR.

The parallel-navigation law is defined by the equations below. Equation 6 guarantees that \vec{r} and $\dot{\vec{r}}$ remain colinear, whereas Equation 7 guarantees that the chaser does not recede from the target. Both equations can be solved for $\dot{\vec{r}}$ in a parametric form using Equation 8 where α is a positive real number. The RG's output is a time-optimal velocity command for the chaser under the parallel-navigation law. As illustrated in Figure 16a, if the chaser maintains a velocity command parallel to the Rendezvous Line (RL), the direction of LOS remains constant, ensuring positional matching between the chaser and the target.

$$\vec{r} \times \dot{\vec{r}} = 0, \quad (6)$$

$$\vec{r} \cdot \dot{\vec{r}} < 0, \quad (7)$$

$$\dot{\vec{r}} = -\alpha \vec{r}. \quad (8)$$

The RG, position matching guidance, is applied to an overtaking problem by creating a virtual target (see Figure 17). Our virtual target is spawned parallel to the predicted trajectory at t_{overtake} point (here, we set 6 seconds). Using the virtual target, we modified the RG to determine whether our autonomy is capable of overtaking or not. Figure 16b visualized the modified RG plot.

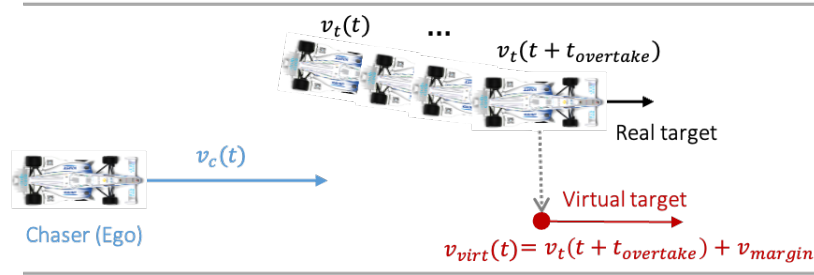


Figure 17. Conceptual visualization of virtual target generation for overtaking decision making.

Considering the maximum acceleration based on the vehicle’s current speed, race flag, and track sector, the Feasible Velocity Region (FVR) is generated. Here, we assumed that the vehicle can accelerate up to 80% of the maximum acceleration for generating the FVR as a safety-performance balance. Finally, the decision module outputs the overtaking trigger signal if the FVR is beyond the Overtaking Rendezvous Line (ORL), which indicates that overtaking is roughly feasible under the current operating configuration. As examples, we visualized two different chaser’s velocity vectors in Figure 16b using v_{c1} and v_{c2} . With the v_{c1} , FVR with blue boxed area is over the ORL, which can be interpreted that our vehicle can be proceed more than the virtual target position at t_{overtake} . On the other hand, v_{c2} with the red boxed area is not long enough to cross the ORL, which means that the ego vehicle will be still behind during the overtaking horizon. As we briefly mentioned before, the size of FVR is decided by multiple factors including vehicle states, race flag, and track sectors. Note that these multiple factors were chosen heuristically.

Minimum Jerkiness Overtaking Trajectory Generation. In this step, the ego-motion trajectory, which will be driven in the near future, is generated. We take the overtaking maneuver as the path switching of the race line and one of the path candidates. Five different reference paths parallel to the track’s center line, $P_{\text{candidate}} \in P_{\text{left}}, P_{\text{left,center}}, P_{\text{center}}, P_{\text{right,center}}, P_{\text{right}}$ were used. Our local trajectory planner is in charge of generating the connecting trajectory segment, $\tau = \{p_{t_0}, p_{t_1}, \dots, p_{t_{ot}}\}$, between the current and the target path, where p_t is the set of position at time t .

Our planner aims to generate a minimum jerkiness trajectory to safely and quickly merge to the target path. Jerkiness is a widely used indicator of comfort in autonomous passenger vehicles, but in the case of a race vehicle, it also impacts the vehicle’s stability. Inspired by (Werling et al., 2010), we generated the minimum jerkiness trajectory by using lateral and longitudinal models in Frenet-frame coordinates as follows:

$$\begin{aligned} d(t) &= c_{d0} + c_{d1}t + c_{d2}t^2 + c_{d3}t^3 + c_{d4}t^4 + c_{d5}t^5, \\ s(t) &= c_{s0} + c_{s1}t + c_{s2}t^2 + c_{s3}t^3 + c_{s4}t^4, \end{aligned} \quad (9)$$

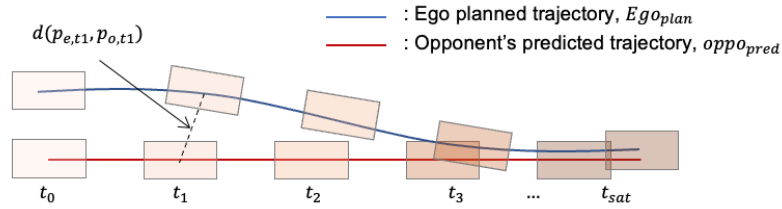
where $c_{i,i=\{d0,d1,d2,d3,d4,d5\}}$ and $c_{i,i=\{s0,s1,s2,s3,s4\}}$ are the coefficients of the lateral and longitudinal components, respectively (we refer to (Werling et al., 2010) for the theoretical proof). Since our goal for the local planner is to merge to the target path, we can set the end condition as $d(t_{ot}) = 0$. Using a set of time, $t_{ot} \in 6, 4, 2$, we generated trajectories with a variety of longitudinal patterns. For the sake of simplicity, we assumed acceleration is zero during planning. By solving Equations 10 and 11, lateral and longitudinal polynomial parameters can be easily calculated given the initial and final states:

$$\begin{bmatrix} t_0^5 & t_0^4 & t_0^3 & t_0^2 & t_0^1 & 1 \\ t_{ot}^5 & t_{ot}^4 & t_{ot}^3 & t_{ot}^2 & t_{ot}^1 & 1 \\ 5t_0^4 & 4t_0^3 & 3t_0^2 & 2t_0^1 & 1 & 0 \\ 5t_{ot}^4 & 4t_{ot}^3 & 3t_{ot}^2 & 2t_{ot}^1 & 1 & 0 \\ 20t_0^3 & 12t_0^2 & 6t_0^1 & 2 & 0 & 0 \\ 20t_{ot}^3 & 12t_{ot}^2 & 6t_{ot}^1 & 2 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} c_{d5} \\ c_{d4} \\ c_{d3} \\ c_{d2} \\ c_{d1} \\ c_{d0} \end{bmatrix} = \begin{bmatrix} d_0 \\ d_{t_{ot}} = 0 \\ \dot{d}_0 \\ \dot{d}_{t_{ot}} = 0 \\ \ddot{d}_0 \\ \ddot{d}_{t_{ot}} = 0 \end{bmatrix}, \quad (10)$$

Algorithm 3. Trajectory clearance checking.

```

1:  $C \leftarrow False$  ▷ Clear collision flag.
2: for  $P_e \in Ego_{plan}$  do
3:   for  $P_o \in Oppo_{pred}$  do
4:      $F \leftarrow SAT(P_e, P_o)$ 
5:     if  $F == True$  then ▷ If any intersection exist.
6:        $t(P_e, P_o) \leftarrow calcTimeDiff(t(P_e, P_o))$ 
7:       if  $t(P_e, P_o) < t_{thres}$  then ▷ Check time difference, here, we set  $t_{thres}$  as 1 sec.
8:         return  $C \leftarrow True$ 
9:       end if
10:    end if
11:  end for
12: end for
13: return  $C$ 
    
```


Figure 18. Collision detection based on opponent's predicted trajectory using SAT algorithm.

$$\begin{bmatrix} t_0^4 & t_0^3 & t_0^2 & t_0^1 & 1 \\ 4t_0^3 & 3t_0^2 & 2t_0^1 & 1 & 0 \\ 4t_{ot}^3 & 3t_{ot}^2 & 2t_{ot}^1 & 1 & 0 \\ 12t_0^2 & 6t_0^1 & 2 & 0 & 0 \\ 12t_{ot}^2 & 6t_{ot}^1 & 2 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} c_{s4} \\ c_{s3} \\ c_{s2} \\ c_{s1} \\ c_{s0} \end{bmatrix} = \begin{bmatrix} s_0 \\ \dot{s}_0 \\ \ddot{s}_0 \\ \dot{s}_1 = 0 \end{bmatrix}. \quad (11)$$

Trajectory Feasibility Checker. The last step of the proposed overtaking planner verifies whether the generated trajectory collides with any predicted trajectory and also is dynamically feasible. For collision checking, a Separating Axis Theorem (SAT) was used (Gottschalk, 1996). Figure 18 and Algorithm 3 illustrate the collision checking procedure visually and algorithmically. For each point in both generated and predicted trajectories, a box shape polygon was created based on the vehicle dimension (see Figure 18), and collisions were detected by checking the intersection between two polygons using SAT. When trajectory clearance checking module confirm that there is no collision, the planned trajectory is investigated whether it also satisfies the tire constraint using the centripetal force equilibrium (Pacejka, 2005). Finally, a collision-free and dynamically feasible trajectory is input to the controller to generate control commands to follow accurately.

5.5. Resilient Control Stack

5.5.1. Overview

The control stack is in charge of generating throttle/brake position and steer angle to follow the planned trajectory accurately. Our high-level control stack is built on various control algorithms for overall system resilience. Control command was calculated strictly at 100 Hz and passed to the drive-by-wire system. Figure 19 visualizes the various types of high-level controllers and their configurations within our control stack.

We noted that MPCC was integrated but only for the simulation race and is not be covered in this paper. For more details of our MPCC controller, refer to (Jung et al., 2021; Liniger et al., 2015).

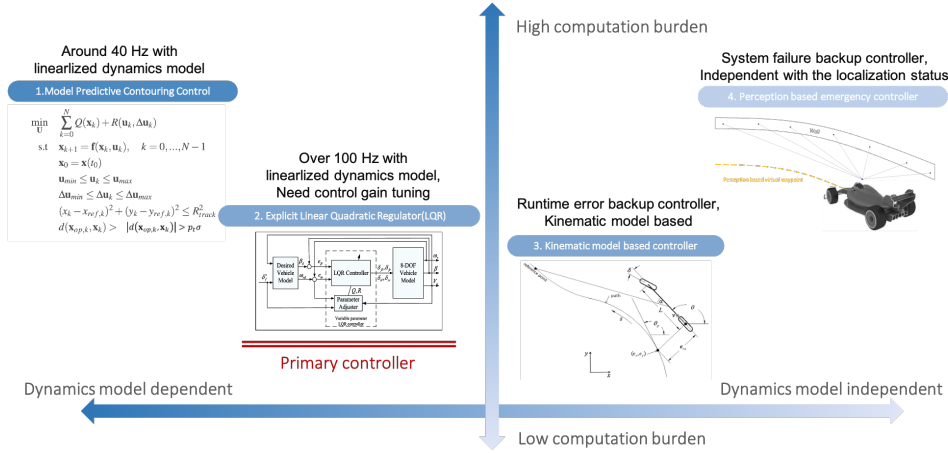


Figure 19. Overview of the control stack, including brief descriptions of each controller. Each controller is assigned a position based on the level of computational burden and model usage. In terms of real-time computability and performance, the LQR-based longitudinal and lateral controller was chosen as the primary controller.

We used a linear–quadratic regulator controller (LQR) as a primary controller for the real-world races. To ensure real-time operability, we created the state feedback gain as a look-up table during the offline optimization. The proposed stack also includes the Stanley controller (Thrun et al., 2006) and the pure pursuit controller (Coulter, 1992). These controllers are based on a simple kinematic model, have a simple implementation, and are exceptionally computationally light. These controllers are used as a backup solution when the primary controller has any issues (e.g., algorithm crash, calculation delay). The solely perception-based emergency controller is another resilience feature of our control stack. As the name suggested, it is independent of the global localization quality and only uses the locally sensed data (here, we used LiDARs points). It is designed to activate when the system status manager reports a fatal error or bad quality from the localization module.

Calculated steer angle and desired acceleration from the high-level controller are passed to the low-level controller to actuate throttle, brake, and steer. Dallara-AV21 is equipped with Schaeffler Paravan SpaceDrive, Drive-by-Wire (DBW) system. For the lateral control, it takes the desired steer angle for the position control. The acceleration command was converted to the pedal positions using engine torque and gear map for the longitudinal control. Implementation details are provided in the following sections.

5.5.2. Explicit LQR-based Controller

As the primary controller of the control stack, we used the LQR-based full-state-feedback controller using linearized system models (Lewis, Vrabie, & Syrmos, 2012; Rajamani, 2011; Spisak et al., 2022).

For the lateral motion model, we used the lateral dynamic bicycle model, whose state elements are represented in terms of position and orientation errors with respect to a reference path. The position error, e_y , and the heading error, e_ψ , are modeled as follows:

$$\frac{d}{dt} \begin{bmatrix} e_y \\ \dot{e}_y \\ e_\psi \\ \dot{e}_\psi \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & -\frac{2C_{\alpha f} + 2C_{\alpha r}}{mV_x} & \frac{2C_{\alpha f} + 2C_{\alpha r}}{m} & -\frac{2C_{\alpha f}l_f - 2C_{\alpha r}l_r}{mV_x} \\ 0 & 0 & 0 & 1 \\ 0 & -\frac{2C_{\alpha f}l_f - 2C_{\alpha r}l_r}{I_z V_x} & \frac{2C_{\alpha f} - 2C_{\alpha r}}{I_z} & -\frac{2C_{\alpha f}l_f^2 + 2C_{\alpha r}l_r^2}{I_z V_x} \end{bmatrix} \begin{bmatrix} e_y \\ \dot{e}_y \\ e_\psi \\ \dot{e}_\psi \end{bmatrix} \tag{12}$$

$$+ \begin{bmatrix} 0 \\ \frac{2C_{\alpha f}}{m} \\ 0 \\ \frac{2C_{\alpha f}l_f}{I_z} \end{bmatrix} \delta + \begin{bmatrix} 0 \\ -\frac{2C_{\alpha f}l_f - 2C_{\alpha r}l_r}{mV_x} - V_x \\ 0 \\ -\frac{2C_{\alpha f}l_f^2 - 2C_{\alpha r}l_r^2}{I_z V_x} \end{bmatrix} \dot{\psi}_{des} + \begin{bmatrix} 0 \\ g \\ 0 \\ 0 \end{bmatrix} \sin(\phi),$$

where $C_{\alpha f}$ and $C_{\alpha r}$ are the cornering stiffness of the front and rear tires, respectively. The model consists of differential equations with respect to the state vector $[e_y, \dot{e}_y, e_\psi, \dot{e}_\psi]^T$ and steering angle control δ . The other terms are governed by the desired yaw rate $\dot{\psi}_{des}$ and bank angle of the track ϕ , which are given by the reference path and track condition. As $\dot{\psi}_{des}$ and ϕ are not included in the state vector, we compensated those terms by using a feedforward control component to minimize the steady-state error (Rajamani, 2011). Without considering the compensated terms, since the mass m , yaw moment of inertia I_z , and distance from the center of gravity to the front and rear axle l_f, l_r are constants, the model can be well-formed as a state space model, assuming the current longitudinal velocity V_x is constant.

Our LQR-based controller's state space model can be then represented as follows:

$$\dot{x}_t = Ax_t + Bu_t, \quad (13)$$

where $x_t = [e_y, \dot{e}_y, e_\psi, \dot{e}_\psi]^T$ and $u_t = \delta$. The optimal state feedback control then is derived as a quadratic programming problem that minimizes a quadratic cost function J as

$$J = \int_0^\infty (x_t^T Q x_t + u_t^T R u_t) dt \quad (14)$$

with gain matrices $Q > 0, R > 0$. The resulting optimal control output $u_t^* = K_t x_t$ is computed with the LQR gain K_t . The original LQR calculates the control gain by solving the algebraic Riccati equation (Lewis et al., 2012). However, the computational cost of solving the equation online is not negligible. Therefore, we derived K_t by solving the Riccati equation offline via the Explicit LQR algorithm (Spisak et al., 2022). The solutions are then used for initializing a set of LQR gains. Therefore, our controller queries a feasible K_t with respect to the current velocity, V_t , without an expensive online optimization process.

The longitudinal model can be described in terms of the nonlinear traction force and aerodynamic drag acting on the ego vehicle. However, since the ego vehicle generates traction force from its engine-based powertrain, it is challenging to model the longitudinal motion with a single dynamic model. Therefore, we designed the longitudinal model to be hierarchical, with a high-level drive-train model and a low-level engine-based powertrain model following the ideas from (Kabzan et al., 2020) and (Rajamani, 2011).

5.5.3. Real-time Perception based Emergency Controller

Most vehicle controllers use global localization and vehicle status, and the performance of localization directly impacts vehicle control performance. A race vehicle driving at high speed in an uncertain localization condition, such as when the GPS sensor responsible for position recognition fails, or the localization result diverges, requires the ability to safely drive (or stop) the vehicle. To this end, our control stack includes a real-time perception-based controller that functions regardless of the current state of localization. As the name implies, the controller employs only real-time environmental sensor data (in this case, LiDARs) to calculate the path the vehicle will take in body coordinates in real-time and generates a control signal to follow. This controller is engaged only when the system manager detects a fatal failure or bad quality of the localization module. The failure modes of localization include GPS signal disconnection or a high estimation of localization uncertainty. The control's conceptual visualization and algorithmic order are depicted in Figure 20.

Our key idea is to generate the virtual local path parallel to the track boundary and safely stop the vehicle without a crash. We used preprocessed LiDAR point cloud (introduced in Section 5.3.2) and conducted polynomial fitting. To minimize fitting error, we chose a model with a smaller fitting error among the first and third-order polynomial models. To minimize the lateral maneuvering in emergency stop scenarios, the virtual path begins from (0,0) in body coordinate if the distance from the wall is over 1.5 m. On the other hand, if the distance to the wall is less than 1.5 m, the virtual path will be generated parallel to the wall with a 1.5 m lateral bias. A pure pursuit controller is in charge of calculating the lateral control.

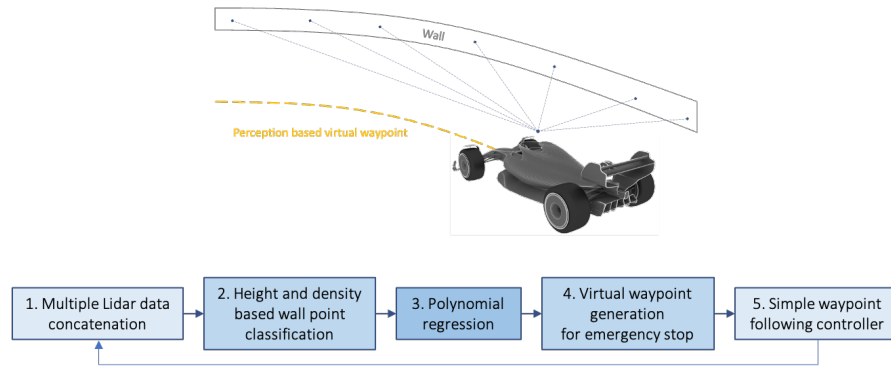


Figure 20. Conceptual visualization and algorithm flow of real-time perception-based lateral vehicle controller.

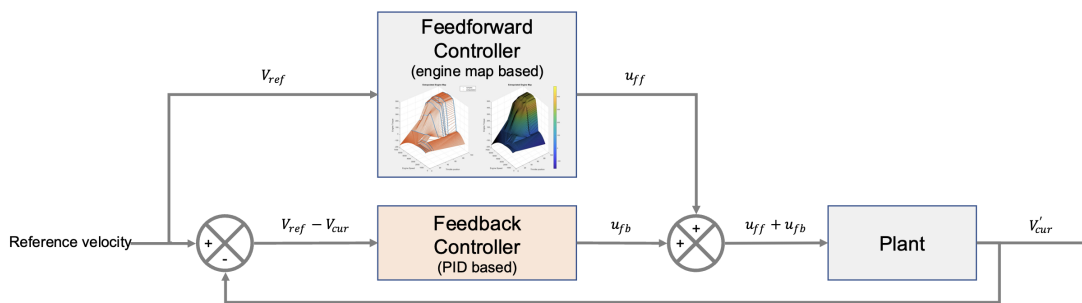


Figure 21. Diagram of a low-level longitudinal controller that outputs the accelerator pedal position from the reference velocity.

5.5.4. Engine Torque Map based Longitudinal Controller

Desired acceleration should finally be converted into throttle and brake positions, inputs of the drive-by-wire module. The longitudinal control diagram is shown in Figure 21. As shown in the diagram, the feedforward control part computes throttle position based on the engine torque map obtained from the chassis dynamometer testing. In designing the feedforward control module, we make the following assumptions to simplify the problem.

1. **Velocity of the vehicle is decided only by engine and brake forces.** Only the throttle and brakes are required to control the vehicle's speed while driving. The track's geometry (e.g., slope, bank) and environmental conditions (e.g., wind) can be ignored.
2. **The torque converter is fully locked.** Torque from the engine passes directly through the transmission without loss.
3. **The tire slip is negligible.** There is no tire slip if the low-level follows the desired velocity.

Deceleration and control errors are corrected using a simple PID-based feedback controller. For more details, please refer to [Seong, Chung, and Shim \(2023\)](#).

6. Experimental Results

6.1. High-speed Solo Lap

This section presents the field testing results of high-speed driving conducted on January 4, 2022, to validate the system's stability before the race event in LVMS. The maximum speed was set to 235 km/h, while the maximum acceleration was set to 3 m/s^2 . The experiment was conducted while

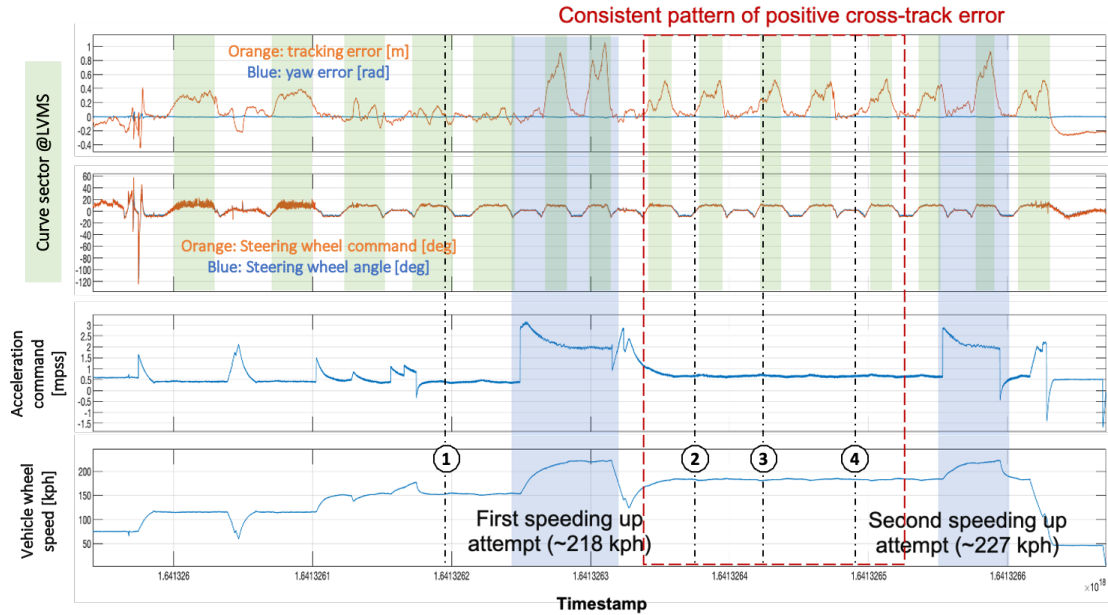


Figure 22. High-speed solo driving experimental results. (First row) Orange and blue lines represent the lateral deviation and yaw error with respect to the reference line. (Second row) Orange and blue lines represent the steer command and position in degree, respectively.

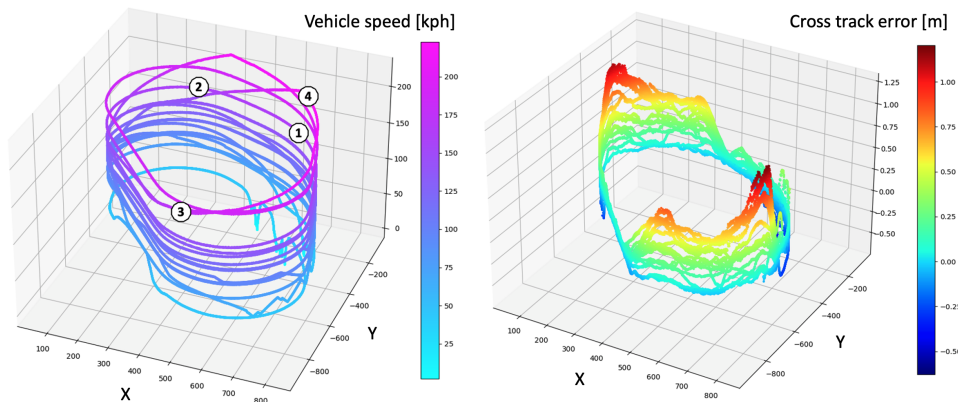


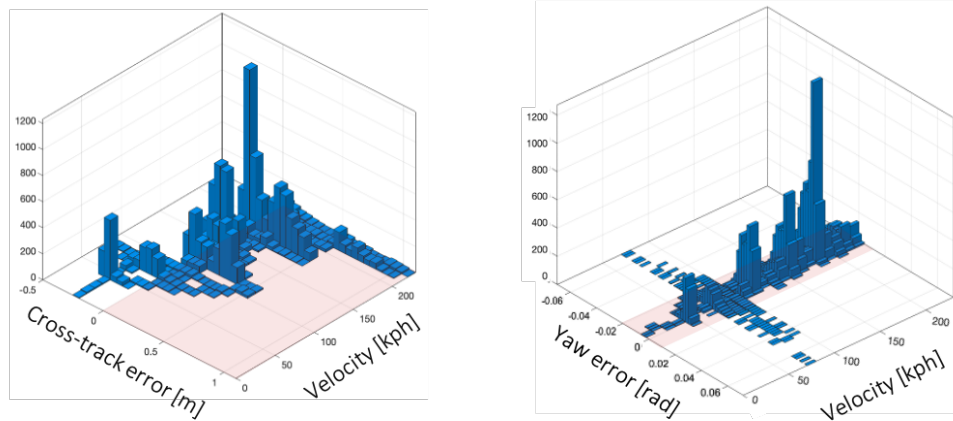
Figure 23. Lateral deviation from the race line plot according to the driven trajectory. The numbers overlaid in the left figure match with timestamps of Figure 22.

gradually increasing the maximum speed. The control result at the time is depicted in Figure 22. In addition, Figure 23 shows the driving speed and lateral deviation.

Figure 24 shows the lateral deviation from the race line. The maximum cross-track error was about 1.15 m at speeds over 220 km/h. Furthermore, 87% of the cross-track errors showed positive signs, which means that our vehicle was on the left side of the race line. This is because the DALLARA-AV21 has an oval setting (left camber setting) which is not modeled in our vehicle model. Also, the yaw error is depicted in Figure 24 on the left. There was a considerable yaw error at 75 km/h because of switching from the pit-out path to the race line on the track. In other cases, 97% of yaw errors are kept within the 0.02 rad. The quantitative analysis results of a high-speed lap are summarized in Table 3. In this experiment, we could drive at a maximum speed of 227.2 km/h.

Table 3. High-speed driving performance analysis result.

Velocity Range [km/h]	Tracking Error [m]			Max Yaw Error [rad]	Max Driving Speed [km/h]
	Max	Absolute Mean	Standard Deviation		
$v < 100$	0.59	0.15	0.192	0.074	
$100 < v < 150$	0.71	0.17	0.214	0.07	227.2
$150 < v$	1.15	0.3	0.443	0.067	

**Figure 24.** Tracking error results. (Left) Lateral deviation. (Right) Yaw error.

6.2. 1:1 Head-to-head Race Event

This section introduces the race results at the IAC’s quarter-final and semifinal rounds. A total of 5 teams that passed qualification participated in the competition: *TUM* (Germany), *KAIST* (South Korea), *PoliMOVE* (Italy), *TII EuroRacing* (Italy), and *Auburn* (USA). The race was held in a tournament format. The team who successfully overtook the opponent advanced to the next round. *PoliMOVE* was given a bye because they demonstrated the fastest driving during qualifications. *Auburn* was our quarter-final rival. Race control set the starting defender’s target velocity at 128 km/h (80 mph). We were instructed by race control to begin the race as an attacker, while *Auburn* started as a defender. To ensure safety, the two vehicles exit the pit box in the order of defender and attacker.

The control result from the pit out in the quarter-final round is shown in Figure 25. We began the race as an attacker attempting to overtake the opposing team. We performed pit-out and formation laps based on race control signals. A green area denotes the section where pit-out was done. During the pit-out mission, the speed was around 50 km/h. As soon as the pitting out is completed, our vehicle aligns with the race line. The formation lap following the race control was shown in the pink area.

Following the race control, the attacker should reduce the gap between himself and the defender. We communicated with the race control during the formation lap and increased the maximum speed. In Figure 25, PB1 indicates when the false detection (hereafter referred as to FD) firstly occurred, and unnecessary deceleration was performed. At the time, our vehicle was running the formation lap following the inner side of the track, and a FD from the perception module occurred near the right track boundary. At the time of PB2, the race control provided an attacker flag, allowing us to attempt overtaking, and our planner switched to the outer line, the attacker mode’s default line. Simultaneously, second phantom braking occurred due to a FD. We increased the attacker mode’s

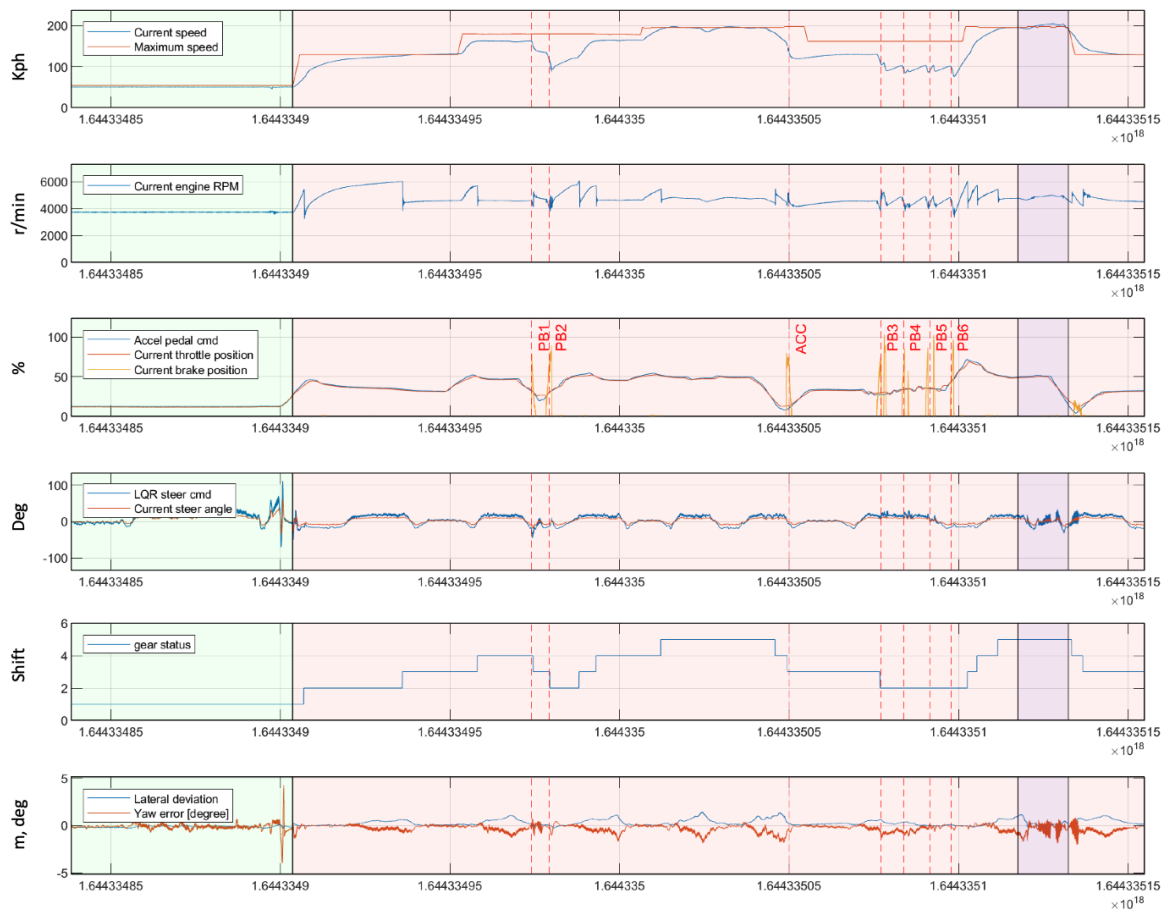


Figure 25. Quarter-final driving result. Green areas indicate pit-outs. The formation lap was performed according to communication with race control, indicated by the pink area. A purple area indicates the section overtaking a vehicle traveling at 128 km/h (80 mph), and the driving speed at that time was about 212 km/h.

speed to 198 km/h to quickly catch the opponent. At ACC, *Auburn*'s vehicle was accurately detected for the first time at a distance of about 100 m and decelerated to the defender speed of 128 km/h (80 mph) to maintain the distance. After a while, four times phantom brakes (PB3, PB4, PB5, and PB6) occurred, and the speed dropped to about 100 km/h (60 mph). After the FD disappeared, our vehicle accelerated again up to our top speed, 200 km/h (124 mph), following the race line. When the distance gap between the two vehicles was about 34 m in the front stretch, our speed was approximately 185 km/h (115 mph), and we started overtaking by safely switching the path to the outer line. After passing, we returned to the inner track line about 32 m of safety margin from *Auburn*'s vehicle. The above-mentioned situations are shown in Figure 26 in order of time. We performed defender-attacker role switching after our overtaking, but the race control halted the race due to an error in the *Auburn*'s system.

As in the quarter-final, we performed a total of three times of defender and two times of attacker in the semifinal with *PoliMOVE*. Figure 27 shows the overall control results in the semifinal. We also experienced phantom braking several times, similar to the quarter-finals. The first time we faced the phantom braking is shown by PB1 in the third row of Figure 27. After the FD disappeared, our vehicle slowed down (ACC1) to maintain the distance while waiting for the front stretch overtaking zone. C1 timestamp shows when the vehicle got the overtaking flag from the race control and switched to the race line since there was no collision within 2 seconds. At EB1, our

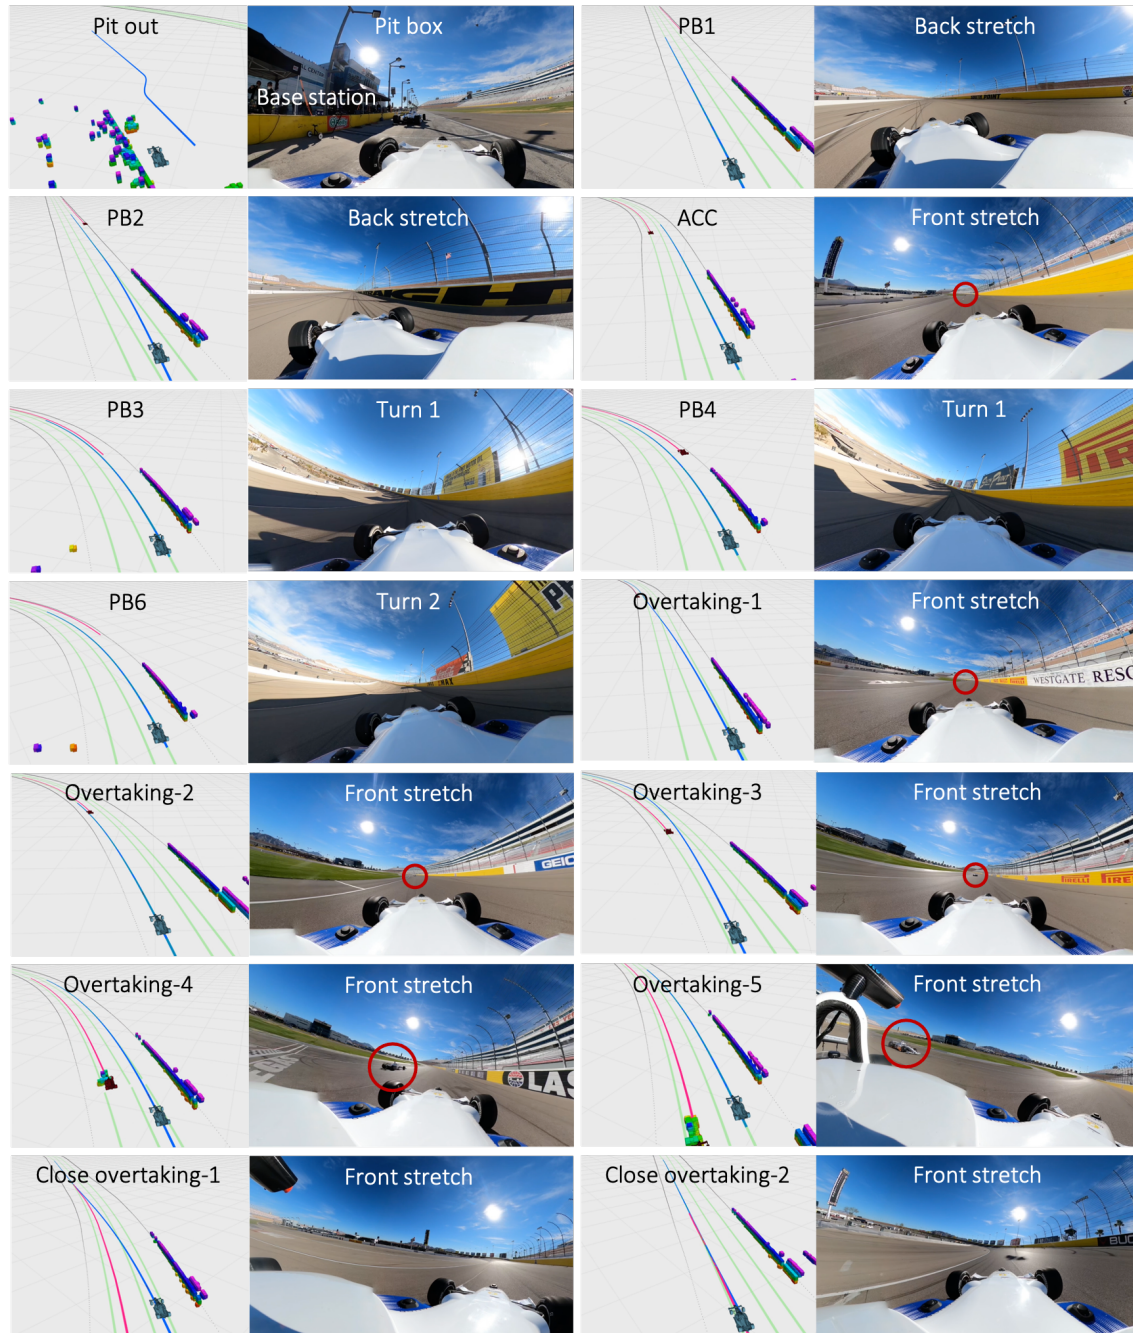


Figure 26. Onboard camera view with perception and planning result visualization results at quarter-final race.

vehicle started to overtake the front vehicle at a speed of 150 km/h (see Figure 28, Overtaking(80)-1, and Overtaking(80)-2). The next round of defender speed was set to 160 km/h (100 mph). PB2 is a phantom braking point that occurred while closing the distance with the *PoliMOVE*, and we accelerated up to 204 km/h (126 mph), successfully overtaking an opposing vehicle traveling at 160 km/h (100 mph). We started overtaking at the end of the front stretch due to phantom braking, and both vehicles were side-by-side at turn 1. Our planner did not permit closing overtaking in the high banking angle zone and completed the overtaking maneuver at the backstretch. However,

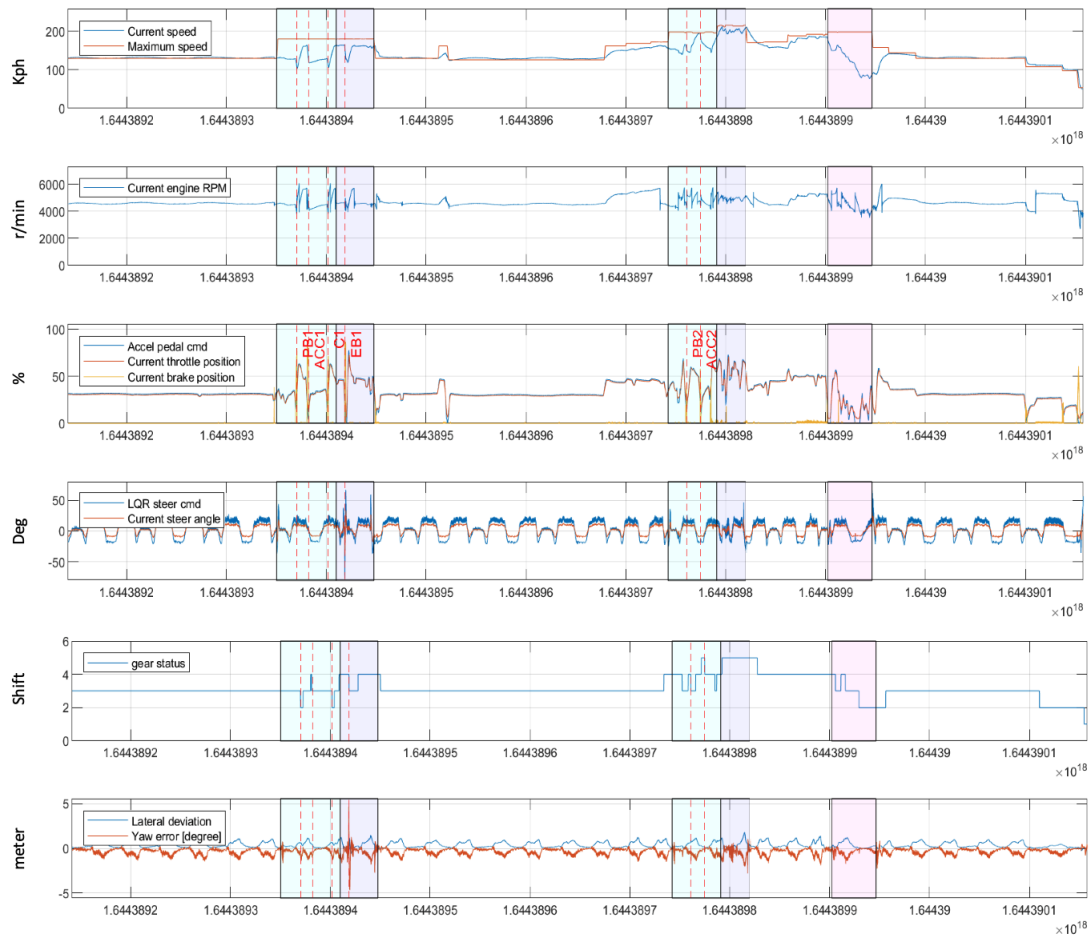


Figure 27. Semifinal driving result. Overtaking in attacker mode (80 and 100 mph) was performed a total of 2 times, and defender role (80, 100, and 115 mph) was performed three times. All missions of pit-in and out, formation lap, and race were performed, and driving was performed for about 18 minutes. When overtaking a vehicle traveling at 100 mph, the top speed was about 204 km/h. FDs caused repeated phantom braking. We decided it was unsafe for both teams, so we finally asked the race control a black flag for a pit-in, and the pink area shows the point in time.

we decided to request a black flag to race control for the retirement because of repeated phantom braking (shown in the pink area of Figure 27).

In the same way as the method introduced in Section 6.1, the control performance in the head-to-head race was analyzed. Figure 29 shows the lateral displacement and yaw error. Figure 29 appears similar to Figure 24, which means that our controller can accurately track the fixed path but also dynamically changing path. The quantitative analysis results are summarized in Table 4. Our maximum speed during the semifinal was 212 km/h (132 mph), and accelerations up to 12.41 m/s^2 . The root cause of all FDs was a mismatch of the right side of the track boundary, especially at the curve sectors. Thus our geofence filter with a 1 m threshold distance sometimes could not correctly remove the LiDAR-based clustering noise or raw radar data.

6.3. Computational Performance Analysis

Real-time systems, relying on either soft or hard timing constraints, require proper diagnostics to verify that the design constraints are respected at run-time. Timing analysis for complex systems is

Table 4. Head-to-head race analysis result. (ABYE: Absolute max yaw error. MLA: Max longitudinal acceleration. AMLA: Absolute max lateral acceleration. MS: Max speed.)

Velocity Range [km/h]	Tracking Error ↓ [m]			ABYE [rad]	MLA [m/s ²]	AMLA [m/s ²]	MS [km/h]
	Max	Absolute Mean	Standard Deviation				
$v < 100$	0.4005	0.1068	0.0995	0.0415			
$100 < v < 150$	1.2318	0.3846	0.2905	0.0815	12.4155 (1.266 g)	6.8750 (0.701 g)	212.5 (132 mph)
$150 < v$	1.7791	0.4872	0.3606	0.0812			



Figure 28. Onboard camera view with perception and planning result visualization results at semifinal race.

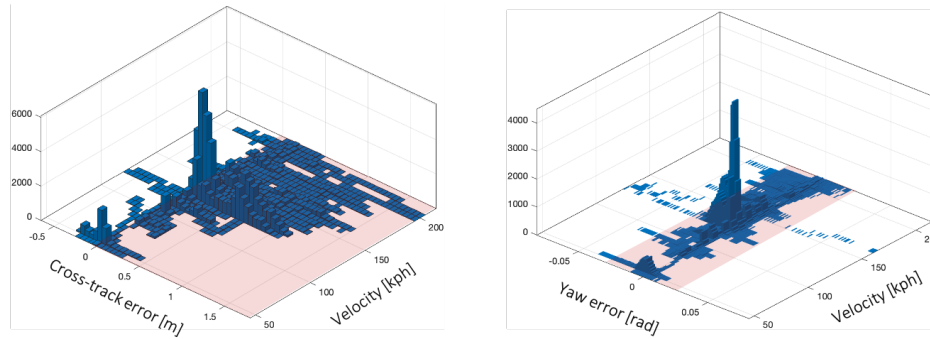


Figure 29. Tracking error histogram according to the driving speed during the CES2022 head-to-head autonomous race. (Left) Lateral deviation. (Right) Yaw error.

Table 5. Timing analysis workstation's specifications.

Component	Specification
CPU	Intel Core i9-20980HK 2.40 GHz (16T, 8C)
GPU	Nvidia GeForce RTX 3080 Laptop x 1
RAM	32 GB
Storage	2 TB Samsung 970Evo NVMe SSD
Software	ROS2 Galactic on Ubuntu 20.04

an active research topic. Recent techniques, including some specifically designed for ROS2 (Bédard, Lütkebohle, & Dagenais, 2022; Li, Hasegawa, & Azumi, 2022), have been released only recently. In this section, an analysis of the timing properties of our proposed autonomy stack is given. The timing measurements have been carried out on the platform detailed in Table 5, by replaying recorded data. The tracing tool adopted to measure the software performance is derived from `Autoware_perf` (Li et al., 2022), which in turn is based on the `ros2_tracing` (Bédard et al., 2022) and `LTTng` (Desnoyers & Dagenais, 2006) tools.

We measured the execution time of the functions that compose the main execution pipeline (registered as ROS2 topic or timer callbacks). For the performance analysis, we plot the number of samples over their execution time (discretizing the timeline). Both the evaluation workstation and the computation system on Dallara-AV21 run on the default Linux 5.8 kernel, which employs a Completely Fair Scheduling algorithm (Pabla, 2009). Figure 30 shows various task duration plots.

The timing requirements set during the design phase (i.e., 100 Hz operation frequency), are met by the stack, according to the data. The nodes composing the localization-planning-control sequence performs well-below the 10 ms deadline, with reasonably low variance (Figures 30a, 30d, and 30e). Similar conclusions can be drawn with regard to the main horizontal component, the System Status Manager (Figure 30f). The perception stack (intended as the detection-tracking-prediction sequence) shows lower performance compatibly with the higher resource-consuming nature of the task. In particular, the lidar clustering algorithm sets its mean execution time at 10.52 ms, peaking at 34.37 ms (Figure 30b), while the geofence filtering function averages at 4.41 ms with a registered maximum of 16.86 ms (Figure 30c). Moreover, while the former's variance is reasonably limited, the spread-out shape of the latter's plot suggests that the function's flow might have an irregular duration. This result, although not optimal, sets the worst-case operation frequency of the perception task around 30 Hz, compatibly with the best-case performance of the available laser-based sensors.

Although the results' scale can be heavily influenced by environmental factors (e.g., hardware and OS configuration, number of concurrent processes, external conditions), the relative proportion of time consumption is well represented by the given data, regardless of the testing environment. As

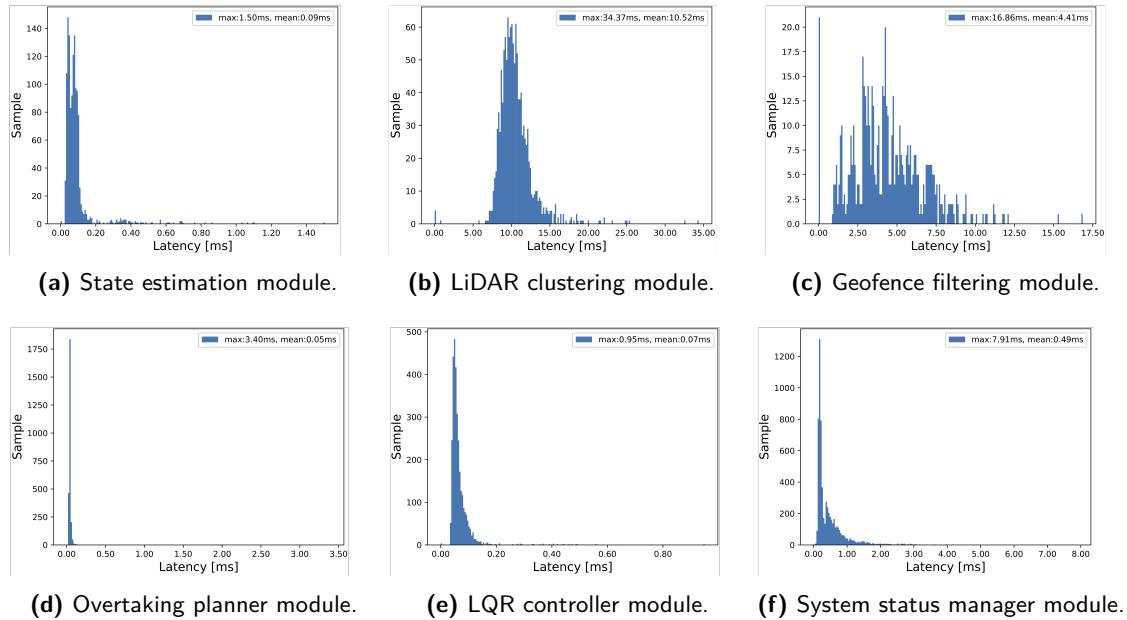


Figure 30. Latency analysis results.

per an evaluation of the end-to-end latency, the data suggest that the average latency sets well-below 10 ms for what concerns the control path (state estimation, planning, control), including the impact of inter-process communication, which averaged around 0.01 ms on our timing analysis workstation.

7. Discussion and Conclusion

In this paper, we presented the full-stack autonomous racing software developed by team KAIST for the Indy Autonomous Challenge (IAC). Our autonomy solution comprises multi-modal perception, a high-speed overtaking planner, a resilient control stack, and a system status manager. All the subsystems of the proposed autonomy stack are developed following our key design principles, aiming to achieve dependability, evolvability, and performance. Even though our autonomy solution is developed targeting the autonomous racing domain, we believe that our system architecture and design principles can be applied to a wide range of robotic applications, especially when it comes to the high-performance, safety-critical, and high-cost-of-failure application domains.

The proposed system was integrated into a full-scaled autonomous race car (Dallara AV-21) and extensively validated through field tests and race events. We, team KAIST, accomplished every mission (including autonomous pit-in/out, static obstacles avoidance, and obeying race flags) in the IAC at Indianapolis Motor Speedway (IMS) in Oct 2021. During IAC at Las Vegas Motor Speedway (LVMS) in Jan 2022, our autonomy demonstrated high-speed head-to-head racing by reaching speeds over 220 km/h and accelerations of up to 12.41 m/s^2 . Our team was one of three teams that successfully finished both race events without system failures or crashes.








Even though we believe that the proposed autonomy solution and the provided results can provide valuable insights to the field robotics community, there are still some technical gaps when it comes to human-like head-to-head racing scenarios. Our autonomy solution was designed based on the IAC race rules, which allowed us to make some fundamental assumptions. For example, our prediction module simplifies the problem by assuming that the other vehicle will maintain a constant speed and a constant displacement from the inner track boundary. However, this assumption is not valid for human-like races. Also, our trajectory planner follows a sampling-based approach which has an advantage in terms of computation burden. However, our planning algorithm has limited capabilities when handling multiple opponents in a competitive scenario. For advanced human-like autonomous

races, a contextual understanding based on learning methods plays a key role. We believe this technology has an enormous impact in multiple application domains ranging from urban autonomous driving to service robots in social environments. At the same time, we would like to leave an open research question: “How can we maintain the system resilience when adopting unobservable learning models?” One promising research topic is explainable deep learning, which has evolved significantly in the last few years. Also, system designs that combine classical and learning-based methods in a complementary way can represent a solid solution to the problem of deployable learning models. With efforts in these fields, we believe that learning algorithms can be more actively adopted in real-world autonomous systems.

Acknowledgments

This work is partially supported by SK Hynix Inc and Institute of Information communications Technology Planning Evaluation (IITP) grant funded by the Korea government (MSIT, 2021-0-00029). We would like to thank Energy System Network (ESN), Juncos Hollinger Racing, and all other participating teams for their support and contributions to the project. Especially, we would like to thank *MIT-PITT-RW* for collaborating on an LQR controller design and camera-based perception module. Also, we thank *PoliMOVE* for sharing their semifinal log data with us, which was used for perception performance evaluation in this work.

ORCID

Chanyoung Jung  <https://orcid.org/0000-0001-8104-2461>
 Hyunki Seong  <https://orcid.org/0000-0002-7169-3006>
 Daegyoo Lee  <https://orcid.org/0000-0002-9336-5759>
 Seungwook Lee  <https://orcid.org/0000-0003-1147-1214>
 Boseong Kim  <https://orcid.org/0000-0002-7437-5169>
 Gyuree Kang  <https://orcid.org/0000-0001-7769-4651>
 Hyunchul Shim  <https://orcid.org/0000-0002-1929-7022>

References

- Ackerman, E. (2022). *Darpa’s racer program sends high-speed autonomous vehicles off-road*. <https://spectrum.ieee.org/darpa-robot-racer>.
- Arya Senna Abdul Rachman, A. (2017). 3d-lidar multi object tracking for autonomous driving: multi-target detection and tracking under urban road uncertainties.
- Bertoncello, M., & Wee, D. (2015). Ten ways autonomous driving could redefine the automotive world. *McKinsey & Company*, 6.
- Betz, J., Zheng, H., Liniger, A., Rosolia, U., Karle, P., Behl, M., ... Mangharam, R. (2022). Autonomous vehicles on the edge: A survey on autonomous vehicle racing. *IEEE Open Journal of Intelligent Transportation Systems*.
- Brodsky, J. S. (2016). Autonomous vehicle regulation: How an uncertain legal landscape may hit the brakes on self-driving cars. *Berkeley Technology Law Journal*, 31(2), 851–878.
- Bédard, C., Lütkebohle, I., & Dagenais, M. (2022). ros2_tracing: Multipurpose low-overhead framework for real-time tracing of ros 2. *IEEE Robotics and Automation Letters*, 7(3), 6511–6518. doi: 10.1109/LRA.2022.3174346
- Cai, P., Mei, X., Tai, L., Sun, Y., & Liu, M. (2020). High-speed autonomous drifting with deep reinforcement learning. *IEEE Robotics and Automation Letters*, 5(2), 1247–1254.
- Carvalho, A., Gao, Y., Gray, A., Tseng, H. E., & Borrelli, F. (2013). Predictive control of an autonomous ground vehicle using an iterative linearization approach. In *16th international ieee conference on intelligent transportation systems (itsc 2013)* (pp. 2335–2340).
- Casanova, D. (2000). On minimum time vehicle manoeuvring: The theoretical optimal lap.
- Christ, F., Wischniewski, A., Heilmeyer, A., & Lohmann, B. (2021). Time-optimal trajectory planning for a race car considering variable tyre-road friction coefficients. *Vehicle system dynamics*, 59(4), 588–612.

- Claussmann, L., Revilloud, M., Gruyer, D., & Glaser, S. (2019). A review of motion planning for highway autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 21(5), 1826–1848.
- Coulter, R. C. (1992). *Implementation of the pure pursuit path tracking algorithm* (Tech. Rep.). Carnegie-Mellon UNIV Pittsburgh PA Robotics INST.
- Dal Bianco, N., Bertolazzi, E., Biral, F., & Massaro, M. (2019). Comparison of direct and indirect methods for minimum lap time optimal control problems. *Vehicle System Dynamics*, 57(5), 665–696.
- De Maesschalck, R., Jouan-Rimbaud, D., & Massart, D. L. (2000). The mahalanobis distance. *Chemometrics and intelligent laboratory systems*, 50(1), 1–18.
- Desnoyers, M., & Dagenais, M. (2006, 01). The ltng tracer: A low impact performance and behavior monitor for gnu/linux. *OLS (Ottawa Linux Symposium)*.
- Elbanhawi, M., & Simic, M. (2014). Sampling-based robot motion planning: A review. *Ieee access*, 2, 56–77.
- Fuchs, F., Song, Y., Kaufmann, E., Scaramuzza, D., & Dürr, P. (2021). Super-human performance in gran turismo sport using deep reinforcement learning. *IEEE Robotics and Automation Letters*, 6(3), 4257–4264.
- Funke, J. (2015). *Collision avoidance up to the handling limits for autonomous vehicles*. Stanford University.
- Funke, J., Brown, M., Erlien, S. M., & Gerdes, J. C. (2016). Collision avoidance and stabilization for autonomous vehicles in emergency scenarios. *IEEE Transactions on Control Systems Technology*, 25(4), 1204–1216.
- Funke, J., Theodosis, P., Hindiyeh, R., Stanek, G., Kritatakirana, K., Gerdes, C., ... Huhnke, B. (2012). Up to the limits: Autonomous audi tts. In *2012 ieee intelligent vehicles symposium* (pp. 541–547).
- Ghumman, U. S., Kunwar, F., Benhabib, B., et al. (2008). *Guidance-based on-line motion planning for autonomous highway overtaking*. University of Toronto.
- Goldfain, B., Drews, P., You, C., Barulic, M., Velev, O., Tsiotras, P., & Rehg, J. M. (2019). Autorally: An open platform for aggressive autonomous driving. *IEEE Control Systems Magazine*, 39(1), 26–55.
- Gottschalk, S. (1996). Separating axis theorem. Technical Report TR96-024. Department of Computer Science. University of North Carolina, Chapel Hill, 20–46.
- Grigorescu, S., Trasnea, B., Cocias, T., & Macesanu, G. (2020). A survey of deep learning techniques for autonomous driving. *Journal of Field Robotics*, 37(3), 362–386.
- Gutman, P.-O., & Velger, M. (1990). Tracking targets using adaptive kalman filtering. *IEEE Transactions on Aerospace and Electronic Systems*, 26(5), 691–699.
- Heilmeier, A., Wischniewski, A., Hermansdorfer, L., Betz, J., Lienkamp, M., & Lohmann, B. (2019). Minimum curvature trajectory planning and control for an autonomous race car. *Vehicle System Dynamics*.
- Herman, J., Francis, J., Ganju, S., Chen, B., Koul, A., Gupta, A., ... Nyberg, E. (2021). Learn-to-race: A multimodal control environment for autonomous racing. In *Proceedings of the ieee/cvf international conference on computer vision* (pp. 9793–9802).
- The history of the indy autonomous challenge powered by cisco*. (2019). <https://www.indyautonomuschallenge.com/history>.
- Howard, A., Sandler, M., Chu, G., Chen, L.-C., Chen, B., Tan, M., ... others (2019). Searching for mobilenetv3. In *Proceedings of the ieee/cvf international conference on computer vision* (pp. 1314–1324).
- Jaritz, M., De Charette, R., Toromanoff, M., Perot, E., & Nashashibi, F. (2018). End-to-end race driving with deep reinforcement learning. In *2018 ieee international conference on robotics and automation (icra)* (pp. 2070–2075).
- Jung, C., Lee, S., Seong, H., Finazzi, A., & Shim, D. H. (2021). Game-theoretic model predictive control with data-driven identification of vehicle model for head-to-head autonomous racing. *arXiv preprint arXiv:2106.04094*.
- Kabzan, J., Valls, M. I., Reijgwart, V. J., Hendriks, H. F., Ehmke, C., Prajapat, M., ... others (2020). Amz driverless: The full autonomous racing system. *Journal of Field Robotics*, 37(7), 1267–1294.
- Karimipour, H., & Dinavahi, V. (2015). Extended kalman filter-based parallel dynamic state estimation. *IEEE transactions on smart grid*, 6(3), 1539–1549.
- Kendall, A., Hawke, J., Janz, D., Mazur, P., Reda, D., Allen, J.-M., ... Shah, A. (2019). Learning to drive in a day. In *2019 international conference on robotics and automation (icra)* (pp. 8248–8254).
- Kunwar, F., & Benhabib, B. (2006). Motion planning for autonomous rendezvous with vehicle convoys. In *2006 ieee intelligent transportation systems conference* (pp. 1568–1573).
- Laumond, J.-P., et al. (1998). *Robot motion planning and control* (Vol. 229). Springer.

- Lee, D., Jung, C., Finazzi, A., Seong, H., & Shim, D. H. (2022). *Resilient navigation and path planning system for high-speed autonomous race car*.
- Lewis, F. L., Vrabie, D., & Syrmos, V. L. (2012). *Optimal control*. John Wiley & Sons.
- Li, Z., Hasegawa, A., & Azumi, T. (2022). Autoware_perf: A tracing and performance analysis framework for ros 2 applications. *Journal of Systems Architecture*, 123, 102341. doi: <https://doi.org/10.1016/j.sysarc.2021.102341>
- Liniger, A., Domahidi, A., & Morari, M. (2015). Optimization-based autonomous racing of 1: 43 scale rc cars. *Optimal Control Applications and Methods*, 36(5), 628–647.
- Liniger, A., & Lygeros, J. (2015). A viability approach for fast recursive feasible finite horizon path planning of autonomous rc cars. In *Proceedings of the 18th international conference on hybrid systems: Computation and control* (pp. 1–10).
- Liniger, A., & Lygeros, J. (2017). Real-time control for autonomous racing based on viability theory. *IEEE Transactions on Control Systems Technology*, 27(2), 464–478.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., & Berg, A. C. (2016). Ssd: Single shot multibox detector. In *European conference on computer vision* (pp. 21–37).
- Lovato, S., & Massaro, M. (2022). A three-dimensional free-trajectory quasi-steady-state optimal-control method for minimum-lap-time of race vehicles. *Vehicle System Dynamics*, 60(5), 1512–1530.
- Madhulatha, T. S. (2012). An overview on clustering methods. *arXiv preprint arXiv:1205.1117*.
- Milliken, W. F., Milliken, D. L., et al. (1995). *Race car vehicle dynamics* (Vol. 400). Society of Automotive Engineers Warrendale, PA.
- O’Kelly, M., Zheng, H., Karthik, D., & Mangharam, R. (2020). Fltenth: An open-source evaluation environment for continuous control and reinforcement learning. *Proceedings of Machine Learning Research*, 123.
- Oshana, R. (2006, 1). Overview of Embedded Systems and Real-Time Systems. *DSP Software Development Techniques for Embedded and Real-Time Systems*, 19–34. doi: 10.1016/B978-075067759-2/50004-1
- Pabla, C. S. (2009). Completely fair scheduler. *Linux Journal*, 2009(184), 4.
- Pacejka, H. (2005). *Tire and vehicle dynamics*. Elsevier.
- Petti, S., & Fraichard, T. (2005). Safe motion planning in dynamic environments. In *2005 ieee/rsj international conference on intelligent robots and systems* (pp. 2210–2215).
- Qin, C., Ye, H., Pranata, C. E., Han, J., Zhang, S., & Liu, M. (2020). Lins: A lidar-inertial state estimator for robust and efficient navigation. In *2020 ieee international conference on robotics and automation (icra)* (pp. 8899–8906).
- Rajamani, R. (2011). *Vehicle dynamics and control*. Springer Science & Business Media.
- Rieber, J. M., Wehlan, H., & Allgower, F. (2004). The roborace contest. *IEEE Control Systems Magazine*, 24(5), 57–60.
- Riedmiller, M., Montemerlo, M., & Dahlkamp, H. (2007). Learning to drive a real car in 20 minutes. In *2007 frontiers in the convergence of bioscience and information technologies* (pp. 645–650).
- Rosolia, U., & Borrelli, F. (2019). Learning how to autonomously race a car: a predictive control approach. *IEEE Transactions on Control Systems Technology*, 28(6), 2713–2719.
- Seong, H., Chung, C., & Shim, D. H. (2023). Data-driven model identification via hyperparameter optimization for the autonomous racing system. *arXiv preprint arXiv:2301.01470*.
- Shan, T., & Englot, B. (2018). Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain. In *2018 ieee/rsj international conference on intelligent robots and systems (iros)* (pp. 4758–4765).
- Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., & Rus, D. (2020). Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping. In *2020 ieee/rsj international conference on intelligent robots and systems (iros)* (pp. 5135–5142).
- Shore, J. (2004, sep). Fail fast. *IEEE Software*, 21(05), 21–25. doi: 10.1109/MS.2004.1331296
- Sommerville, I. (2015). *Software engineering* (10th ed.). Boston: Pearson.
- Spisak, J., Saba, A., Suvarna, N., Mao, B., Zhang, C. T., Chang, C., ... Ramanan, D. (2022). Robust modeling and controls for racing on the edge. *arXiv preprint arXiv:2205.10841*.
- Stahl, T., Wischniewski, A., Betz, J., & Lienkamp, M. (2019). Multilayer graph-based trajectory planning for race vehicles in dynamic scenarios. In *2019 ieee intelligent transportation systems conference (itsc)* (pp. 3149–3154).
- Subosits, J. K., & Gerdes, J. C. (2019). From the racetrack to the road: Real-time trajectory replanning for autonomous driving. *IEEE Transactions on Intelligent Vehicles*, 4(2), 309–320.

- Talvala, K. L., Kritayakirana, K., & Gerdes, J. C. (2011). Pushing the limits: From lanekeeping to autonomous racing. *Annual Reviews in Control*, 35(1), 137–148.
- Thrun, S. (2002). Probabilistic robotics. *Communications of the ACM*, 45(3), 52–57.
- Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., ... others (2006). Stanley: The robot that won the darpa grand challenge. *Journal of field Robotics*, 23(9), 661–692.
- Timings, J. P., & Cole, D. J. (2013). Minimum maneuver time calculation using convex optimization. *Journal of Dynamic Systems, Measurement, and Control*, 135(3), 031015.
- Uppada, S. K. (2014). Centroid based clustering algorithms—a clarion study. *International Journal of Computer Science and Information Technologies*, 5(6), 7309–7313.
- Wan, E. A., Van Der Merwe, R., & Haykin, S. (2001). The unscented kalman filter. *Kalman filtering and neural networks*, 5(2007), 221–280.
- Weiss, T., & Behl, M. (2020). Deepracing: a framework for autonomous racing. In *2020 design, automation & test in europe conference & exhibition (date)* (pp. 1163–1168).
- Werling, M., Ziegler, J., Kammel, S., & Thrun, S. (2010). Optimal trajectory generation for dynamic street scenarios in a frenet frame. In *2010 ieee international conference on robotics and automation* (pp. 987–993).
- Williams, G., Drews, P., Goldfain, B., Rehg, J. M., & Theodorou, E. A. (2016). Aggressive driving with model predictive path integral control. In *2016 ieee international conference on robotics and automation (icra)* (pp. 1433–1440).
- Wurman, P. R., Barrett, S., Kawamoto, K., MacGlashan, J., Subramanian, K., Walsh, T. J., ... others (2022). Outracing champion gran turismo drivers with deep reinforcement learning. *Nature*, 602(7896), 223–228.

How to cite this article: Jung, C., Finazzi, A., Seong, H., Lee, D., Lee, S., Kim, B., Kang, G., & Shim, H. (2023). An autonomous racing system: Design, implementation, and analysis; team KAIST at the IAC. *Field Robotics*, 3, 766–800.

Publisher’s Note: Field Robotics does not accept any legal responsibility for errors, omissions or claims and does not provide any warranty, express or implied, with respect to information published in this article.