Real-Time Distributed Model Predictive Control with Gaussian Process Learning for Planar Vehicle Platooning

Lorenzo LAGANÀ[1], Andrea WRONA[1], Emanuele DE SANTIS[1,*]

Abstract— **This paper proposes a distributed Model Predictive Control framework for planar vehicle platooning that integrates Gaussian Process Regression to compensate unmodeled dynamics and propagate uncertainty within the prediction horizon. The controller is developed consistently with a communication-aware supervisory philosophy and assessed under representative operating conditions associated with varying vehicle-to-vehicle communication quality. It is implemented as a real-time sparse sequential quadratic programming scheme and evaluated in multi-vehicle simulations, including emergency braking, lane changes, and communication-loss scenarios. Results show that a single controller core can ensure safe spacing, coordinated motion, and robust path tracking across diverse platooning conditions.**

I. INTRODUCTION

Intelligent Transportation Systems (ITS) aim to improve road safety, traffic efficiency, and environmental sustainability [1]. Within this framework, vehicle platooning has emerged as a prominent cooperative-driving strategy in which vehicles travel in tightly coordinated formations. Platooning offers benefits such as reduced aerodynamic drag, increased roadway capacity, and enhanced safety through anticipatory control [2], [3]. These advantages, however, rely on strong inter-vehicle coupling and dependable vehicle-to-vehicle (V2V) communication.

Most existing platooning approaches adopt simplified longitudinal vehicle models to enable tractable string-stability analysis and controller design [4]. While effective for formal analysis, such abstractions limit applicability in realistic scenarios involving simultaneous longitudinal and lateral maneuvers, including emergency braking, lane-change avoidance, and curved-road driving [5]. In these conditions, vehicle motion is inherently two-dimensional, and the coupling between longitudinal and lateral dynamics becomes significant, particularly near handling limits.

From a control perspective, a variety of approaches have been proposed, including classical Adaptive Cruise Control (ACC) and Cooperative Adaptive Cruise Control (CACC) strategies [6], [7], consensus-based distributed control laws [8], and more recently, distributed Model Predictive Control (MPC) frameworks [9]. Despite their differences, most of these methods remain rooted in one-dimensional formulations, where lateral control is either neglected or treated as a separate low-level problem. While this separation is acceptable under nominal conditions, it becomes restrictive in safety-critical scenarios requiring coordinated planar motion and constraint handling.

In parallel, communication-aware platooning has gained increasing attention, recognizing that V2V communication quality directly impacts achievable performance, stability, and safety [10]. Several works have proposed adaptive topologies, variable spacing policies, and supervisory strategies to cope with communication degradation [11]. In particular, Graceful Degradation and Upgradation (GDU) mechanisms allow platoons to switch between cooperative and fallback modes under intermittent communication [12], [13]. Nevertheless, these approaches are typically built on simplified longitudinal controllers and are rarely integrated within a unified predictive control framework that explicitly accounts for planar dynamics and constraints.

MPC is a natural candidate for addressing these challenges, as it enables systematic handling of multi-variable dynamics and constraints while providing a predictive framework for coordination [9]. At the same time, learning-based techniques such as Gaussian Process Regression (GPR) have shown strong potential in capturing unmodeled dynamics and improving prediction accuracy [14]–[16]. The integration of GPR within MPC has demonstrated excellent performance in high-performance autonomous driving applications, particularly in scenarios where vehicles operate close to physical limits. However, these methods have been predominantly developed for single-vehicle settings and do not directly address distributed platooning, inter-vehicle coupling, or communication-aware operation.

A key gap in the literature is the lack of a unified and implementable framework combining planar vehicle dynamics, distributed predictive control, learning-based model enhancement, and communication-aware supervision. While recent work has moved in this direction [17], it relies on fully nonlinear MPC formulations that are difficult to implement in real time and do not explicitly account for varying communication modes, including standard ACC operation. To address these limitations, this paper proposes a communication-aware distributed control architecture that integrates nonlinear MPC, Gaussian Process residual modeling, and operating regimes inspired by the GDU framework of Hasan et al. [12]. In contrast to [17], real-time feasibility is achieved through a sparse Sequential Quadratic Programming (SQP) formulation inspired by [14], combined with tractable uncertainty propagation and constraint tightening techniques as in [15]. The resulting framework seamlessly accommodates cooperative, degraded, and ACC operating

*Corresponding author.

[1]Lorenzo Laganà, Andrea Wrona and Emanuele De Santis are with Department of Computer, Control and Management Engineering "Antonio Ruberti", Sapienza University of Rome, Rome, Italy lagana.1937829@studenti.uniroma1.it, {wrona, edesantis}@diag.uniroma1.it

regimes within a single controller.

The main contributions of this work are fourfold: (i) the development of a distributed nonlinear MPC framework for cooperative planar platooning; (ii) the integration of Gaussian Process residual modeling within the prediction layer to capture unmodeled dynamics and propagate uncertainty; (iii) the design and validation of a single controller core across communication-dependent operating regimes inspired by the GDU framework of Hasan et al. [12]; (iv) a real-time implementable GPR-SQP-MPCC formulation, validated in multi-vehicle simulations.

The remainder of this paper is organized as follows. Section II presents the system model and the considered V2V communication operating conditions. Section III details the proposed methodology, including the GP-augmented prediction model, uncertainty propagation, and the GPR-SQP-MPCC formulation. Section IV discusses the simulation results. Finally, Section V concludes the paper.

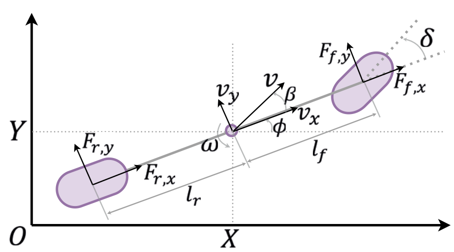

Fig. 1: Single-track bicycle planar model and conventions.

II. System model

The system model combines a nominal planar vehicle model with a data-driven residual term. The nominal dynamics is described by the nonlinear single-track bicycle model (1), based on the scheme in Fig. 1.

$$\begin{bmatrix} \dot{X} \\ \dot{Y} \\ \dot{\phi} \\ \dot{v}_x \\ \dot{v}_y \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} v_x \cos\phi - v_y \sin\phi \\ v_x \sin\phi + v_y \cos\phi \\ \omega \\ \frac{1}{m}\left(F_{r,x} + F_{f,x}\cos\delta - F_{f,y}\sin\delta + m v_y \omega\right) \\ \frac{1}{m}\left(F_{r,y} + F_{f,x}\sin\delta + F_{f,y}\cos\delta - m v_x \omega\right) \\ \frac{1}{I_z}\left(F_{f,x} l_f \sin\delta + F_{f,y} l_f \cos\delta - F_{r,y} l_r\right) \end{bmatrix} \tag{1}$$

The state is $x = [X, Y, \phi, v_x, v_y, \omega]^\top$, where (X, Y) are the Center of Gravity (CoG) coordinates in the inertial frame, ϕ is the yaw angle, (v_x, v_y) are the body-frame longitudinal and lateral velocities, and ω is the yaw rate. The input is $u = [D, \delta]^\top$, where $D \in [-1, 1]$ is the normalized rear-wheel longitudinal actuation and δ is the front steering angle. The parameters m, I_z, and (l_f, l_r) denote the vehicle mass, yaw inertia, and distances from the CoG to the front and rear axles, respectively. The model assumes planar motion, front-wheel steering only, and an algebraic longitudinal actuation law. The nominal predictor used in the controller assumes the lateral tire forces as linear functions of the slip angles $F_{f,y} = C_f \alpha_f, F_{r,y} = C_r \alpha_r$, with $\alpha_f = -\arctan\left(\frac{\omega l_f + v_y}{v_x}\right) + \delta$, $\alpha_r = \arctan\left(\frac{\omega l_r - v_y}{v_x}\right)$ where C_f and C_r are constant cornering stiffnesses. The rear longitudinal tire force is modeled as

$$F_{r,x} = \left(C_{m1} - C_{m2} v_x\right) D - C_{\text{roll}} - C_d v_x^2,$$

where C_{m1}, C_{m2} describe the lumped motor characteristic, C_{roll} the rolling resistance, and C_d the aerodynamic drag. Allowing $D < 0$ enables the same algebraic law to represent braking as well. In simulation, (1) is used as the prediction/control model, while the high-fidelity plant is given by the Vehicle Body 3DOF block from the Vehicle Dynamics Blockset [18], which captures load transfer, aerodynamic effects, and nonlinear lateral tire dynamics consistently with the autonomous-racing benchmark in [14].

A. V2V Communication Model

The V2V layer is modeled by a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, A)$, where $a_{ij} = 1$ if vehicle i receives information from vehicle j, and $\mathcal{N}_i$ is the corresponding neighbour set.

At each sampling instant, vehicle i broadcasts the first two moments of its predicted planar position. The receiving vehicles then reconstruct the corresponding curvilinear progress along the common centerline from the communicated planar information. These neighbour predictions are treated as known parameters in the SQP and enter the controller through the cooperative progress reference. The reconstructed preceding-vehicle curvilinear estimate is used in the rear-end collision-avoidance constraint, while the uncertainty terms are used for chance-based tightening.

The cooperative longitudinal target is generated from the neighbour-average law in [17]:

$$s_i^{\text{ref}}(k) = \sum_{j \in \mathcal{N}_i} n_{ij} \Big(s_j(k) + (j - i) d_0 \Big), \tag{2}$$

with $\sum_{j \in \mathcal{N}_i} n_{ij} = 1,\ n_{ij} \geq 0,\ n_{ii} = 0$ and d_0 being the adopted spacing offset. Thus, all vehicles share the same geometric path, while the along-path reference is adjusted according to the communication topology and the received neighbour trajectories.

III. Proposed Methodology

This work proposes a Model Predictive Contouring Control (MPCC) problem for vehicle platooning, as it combines path tracking and speed planning within a single receding-horizon problem [14]. Moreover, to account for unmodeled dynamics and disturbances, GPR is used to learn the residual mismatch between the nominal predictor and the high-fidelity plant from previously collected state–input data [19], [20]. In the following, the superscript i denotes the vehicle index and $t = 0, \ldots, N_p - 1$ the prediction-step index. For vehicle i, the Gaussian Process (GP)-augmented predictor is used along the horizon with

$$x_{t+1}^i = f_{nom}^i(x_t^i, u_t^i) + B_d\big(d^i(z_t^i) + \nu_t^i\big), \tag{3}$$

where z_t^i is defined as

$$z_t^i = \begin{bmatrix} v_{x,t}^i & v_{y,t}^i & \omega_t^i & D_t^i & \delta_t^i \end{bmatrix}^\top. \tag{4}$$

Since the nonlinear dynamics and the stochastic GP correction make the exact multi-step predictive distribution intractable, a Gaussian closure is adopted and only the first two moments are propagated [20]. Define $w_t^i := \begin{bmatrix}(x_t^i)^\top & (u_t^i)^\top\end{bmatrix}^\top \sim \mathcal{N}(\mu_{w,t}^i, \Sigma_{w,t}^i)$, with

$$\mu_{w,t}^i = \begin{bmatrix}(\mu_{x,t}^i)^\top \, (\mu_{u,t}^i)^\top\end{bmatrix}^\top, \, \Sigma_{w,t}^i = \begin{bmatrix}\Sigma_{x,t}^i & \Sigma_{xu,t}^i \\ (\Sigma_{xu,t}^i)^\top & \Sigma_{u,t}^i\end{bmatrix}$$

Since the regressor is affine in (x_t^i, u_t^i),

$$z_t^i = S\, w_t^i, \qquad S = \begin{bmatrix}B_d^\top & 0 \\ 0 & I\end{bmatrix}, \qquad \bar{z}_t^i = S\, \mu_{w,t}^i.$$

The predictive mean is propagated by evaluating the GP-augmented predictor at the current mean pair,

$$\mu_{x,t+1}^i = f_{nom}^i(\mu_{x,t}^i, \mu_{u,t}^i) + B_d\, \mu_{GP}^i(\bar{z}_t^i). \tag{5}$$

To derive a tractable covariance recursion, the predictor is linearized around the current mean trajectory. Let $J_t^i := \partial_z \mu_{GP}^i\big|_{\bar{z}_t^i}$. and define

$$A_{x,t}^i := \left.\frac{\partial f_{nom}^i}{\partial x}\right|_{(\mu_{x,t}^i, \mu_{u,t}^i)} + B_d J_t^i \begin{bmatrix}B_d^\top & 0\end{bmatrix}, \tag{6}$$

$$A_{u,t}^i := \left.\frac{\partial f_{nom}^i}{\partial u}\right|_{(\mu_{x,t}^i, \mu_{u,t}^i)} + B_d J_t^i \begin{bmatrix}0 & I\end{bmatrix}. \tag{7}$$

Under a first-order approximation, the state covariance evolves as

$$\begin{aligned}\Sigma_{x,t+1}^i \approx\ & A_{x,t}^i \Sigma_{x,t}^i (A_{x,t}^i)^\top + A_{x,t}^i \Sigma_{xu,t}^i (A_{u,t}^i)^\top \\ & + A_{u,t}^i (\Sigma_{xu,t}^i)^\top (A_{x,t}^i)^\top + A_{u,t}^i \Sigma_{u,t}^i (A_{u,t}^i)^\top \\ & + B_d\big(\Sigma_d^{GP,i}(\bar{z}_t^i) + \Sigma_\nu\big) B_d^\top.\end{aligned} \tag{8}$$

To limit covariance growth, an ancillary linear feedback is introduced: $u_t^i = \mu_{u,t}^i + K^i(\mu_{x,t}^i - x_t^i)$, which implies $\Sigma_{u,t}^i = K^i \Sigma_{x,t}^i (K^i)^\top$ and $\Sigma_{xu,t}^i = -\Sigma_{x,t}^i (K^i)^\top$. Therefore, (8) reduces to

$$\Sigma_{x,t+1}^i \approx A_{cl,t}^i \Sigma_{x,t}^i (A_{cl,t}^i)^\top + B_d\big(\Sigma_d^{GP,i}(\bar{z}_t^i) + \Sigma_\nu\big) B_d^\top,$$

with $A_{cl,t}^i := A_{x,t}^i - A_{u,t}^i K^i$. Hence, the GP posterior mean corrects the nominal prediction through (5), while the posterior covariance enters the MPC through uncertainty propagation and the associated chance-constraint tightening.

A. GPR-SQP-MPCC Formulation

The implemented controller is a real-time sparse reformulation of the nonlinear distributed stochastic GP-based MPC framework in [17], combining the GP-oriented approximation strategy of [20] with the MPCC structure of [14] and the constraint treatment of [15].

Rather than solving a fully coupled stochastic optimal control problem online, the controller adopts an SQP approximation in which the GP posterior mean is evaluated and linearized along the shifted scheduling trajectory, so that its first-order effect is embedded in the local SQP dynamics (6)–(7). In contrast, covariance propagation is performed externally along the same scheduling trajectory and is used only to update chance-constraint tightenings before the SQP solve. As a result, the online SQP optimizes the predicted mean dynamics under pre-tightened constraints, while remaining blind to decision-dependent changes in covariance within the current iteration. The resulting online controller solves a sequence of sparse QPs in receding-horizon fashion.

The prediction model is augmented with an explicit curvilinear progress state s and a virtual progress-rate input v_θ. The augmented state and input are defined as

$$x_k := \begin{bmatrix}X_k & Y_k & \psi_k & v_{x,k} & v_{y,k} & \omega_k & s_k\end{bmatrix}^\top, \tag{9}$$

$$u_k := \begin{bmatrix}D_k & \delta_k & v_{\theta,k}\end{bmatrix}^\top, \tag{10}$$

where the physical inputs are the normalized longitudinal actuation D and the front steering angle δ, while v_θ governs the progress dynamics $s_{k+1} = s_k + T_s v_{\theta,k}, \;\; 0 \le v_{\theta,k} \le \bar{v}_\theta$.

The nonlinear MPCC problem is solved by successive convexification around a shifted warm-start trajectory. The initial trajectory is obtained by shifting the previous solution by one stage, injecting the current state at the first node, and completing the tail with the last available input. At each SQP pass, all nonlinear quantities are re-evaluated along the current trajectory, the GP-corrected dynamics and nonlinear terms are linearized, and the resulting sparse multiple-shooting QP is solved. The QP solution then defines the next linearization trajectory. After the last pass, only the first control move is applied to the plant.

Let $g_t(x_t, u_t)$ denote the GP-corrected discrete-time prediction map. Linearizing around the current trajectory point $(\bar{x}_t, \bar{u}_t)$ gives $g_t(x_t, u_t) \approx g_t(\bar{x}_t, \bar{u}_t) + A_t(x_t - \bar{x}_t) + B_t(u_t - \bar{u}_t)$, with $A_t = \partial_x g_t\big|_{(\bar{x}_t, \bar{u}_t)}$, $B_t = \partial_u g_t\big|_{(\bar{x}_t, \bar{u}_t)}$. Equivalently, the predictor used in the QP is written as

$$x_{t+1} \approx A_t x_t + B_t u_t + c_t, \qquad c_t = g_t(\bar{x}_t, \bar{u}_t) - A_t \bar{x}_t - B_t \bar{u}_t,$$

which matches the nonlinear GP-corrected model at the expansion point and yields the stage-wise affine form required by sparse multiple shooting.

To improve numerical conditioning, the implemented SQP is formulated in scaled coordinates using diagonal scaling matrices T_x and T_u and $T_{\Delta u}$. To impose input-rate bounds linearly, the decision vector stacks all stage-wise state and input variables, together with the input increments $\Delta u_t := u_{t+1} - u_t$. The resulting QP has the standard sparse form

$$\min_{Z_k} \frac{1}{2} Z_k^\top H Z_k + f^\top Z_k \tag{11}$$

subject to

$$A_{\mathrm{eq}} Z_k = b_{\mathrm{eq}}, \qquad l_g \le G Z_k \le u_g, \qquad l_b \le Z_k \le u_b.$$

The stage objective retains the MPCC structure while adding a cooperative longitudinal coordination term based on the neighbour-average reference trajectory (2). Define $e_t := \begin{bmatrix}\hat{e}_c(x_t, s_t) & \hat{e}_l(x_t, s_t)\end{bmatrix}^\top$, $r_t := \begin{bmatrix}\omega_t & e_\phi(x_t, s_t)\end{bmatrix}^\top$, $Q_r := \mathrm{diag}(q_\Omega, q_\phi)$, with $e_\phi(x_t, s_t) := \psi_t - \Phi(s_t)$. The stage cost is then written as

$$\ell_t = \|e_t\|_{Q_e}^2 + \|u_t\|_R^2 + \|\Delta u_t\|_S^2 + q_S(s_t - s_t^{\mathrm{ref}})^2 + \|r_t\|_{Q_r}^2.$$

Here, $\|e_t\|_{Q_e}^2$ penalizes contouring and lag errors with respect to the common centerline-based path, while $q_S(s_t -$

$s_t^{\mathrm{ref}})^2$ drives the vehicle toward the cooperative along-path target. The terms $\|u_t\|_R^2$ and $\|\Delta u_t\|_S^2$ penalize the control effort and its stage-to-stage variation, respectively, the latter promoting smoother inputs and discouraging aggressive actuation changes. Finally, $\|r_t\|_{Q_r}^2$ provides weak regularization on yaw rate and heading misalignment, while the dominant geometric tracking action is provided by contouring and lag.

To obtain the quadratic stage cost used in the SQP, the nonlinear MPCC error map is linearized at first order around the current SQP trajectory:

$$e_t(x_t) \approx \bar{e}_t + G_t(x_t - \bar{x}_t), \quad \bar{e}_t := e_t(\bar{x}_t), \quad G_t := \left.\frac{\partial e_t}{\partial x}\right|_{\bar{x}_t}.$$

Equivalently, $e_t(x_t) \approx G_t x_t + d_t^e$, and $d_t^e := \bar{e}_t - G_t \bar{x}_t$. Substituting this approximation into the weighted tracking term yields $\|e_t\|_{Q_e}^2 \approx x_t^\top Q_{e,t} x_t + q_{e,t}^\top x_t + c_{e,t}$, with $Q_{e,t} := G_t^\top Q_e G_t,\ q_{e,t} := 2G_t^\top Q_e d_t^e$.

The heading-error contribution is treated analogously through a first-order linearization of $e_\phi(x_t, s_t)$, whereas the cooperative progress, yaw-rate, absolute-input, and input-increment penalties are already quadratic and therefore enter the QP directly through selector-based state blocks and diagonal input blocks.

Defining the stage variable $\chi_t := \begin{bmatrix} x_t^\top & u_t^\top & \Delta u_t^\top \end{bmatrix}^\top$, the stage contribution can be written as

$$\ell_t^{\mathrm{QP}} = \frac{1}{2}\chi_t^\top H_t \chi_t + h_t^\top \chi_t + \mathrm{const}, \tag{12}$$

with $H_t = 2\,\mathrm{blkdiag}(Q_{x,t},\ R,\ S)$, $h_t = \begin{bmatrix} q_{e,s,\phi,t} & 0 & 0 \end{bmatrix}^\top$, where $Q_{x,t}$ and $q_{e,s,\phi,t} := q_{e,t} + q_{s,t} + q_{\phi,t}$ collect the state-side contributions of MPCC tracking, cooperative progress, yaw-rate and heading regularization. Terminal weights are handled by replacing the last-stage tracking blocks with heavily amplified terminal counterparts, thereby mimicking terminal constraints without shrinking the feasible set.

B. Problem Constraints

The implemented controller includes the following classes of constraints.

1) Scaled Box and Rate Constraints: State and input bounds are imposed directly in the scaled QP coordinates: $x_{\min} \le x_t \le x_{\max}$, $u_{\min} \le u_t \le u_{\max}$, together with $\Delta u_{\min} \le \Delta u_t \le \Delta u_{\max}$. These bounds are enforced deterministically on the optimized mean trajectory.

2) Lane-Keeping via Time-Varying Slab Constraints: Following [14] with covariance-aware tightening of [15], lane keeping is enforced through time-varying slab constraints. At each stage t, the admissible lane region is constructed from the known left and right border endpoints $p_{L,t}, p_{R,t} \in \mathbb{R}^2$. Defining $d_t := p_{R,t} - p_{L,t}$, $a_t := -d_t$, the deterministic slab bounds are obtained as

$$db_{1,t} := a_t^\top p_{L,t}, \qquad db_{2,t} := a_t^\top p_{R,t},$$

$$db_{\max,t} := \max(db_{1,t}, db_{2,t}), \quad db_{\min,t} := \min(db_{1,t}, db_{2,t}),$$

so that the nominal lane constraint reads

$$db_{\min,t} \le a_t^\top p_t \le db_{\max,t}, \qquad p_t := \begin{bmatrix} X_t & Y_t \end{bmatrix}^\top.$$

Considering the uncertainty-aware tightening, the predicted planar covariance Σ_t^{xy} is projected onto the transverse direction

$$\hat{n}_t := \frac{d_t}{\|d_t\|}, \qquad \sigma_{n,t} := \sqrt{\hat{n}_t^\top \Sigma_t^{xy} \hat{n}_t},$$

and the corresponding lateral margin is computed as $\Delta_{\mathrm{lat},t} := \min\big(\beta_{\mathrm{lane}}\sigma_{n,t}, \Delta_{\mathrm{lat}}^{\max}\big)$. This margin is then converted into slab-coordinate units through $\Delta_{db,t} := \|d_t\|\,\Delta_{\mathrm{lat},t}$, yielding the tightened bounds $ug_{\mathrm{lane},t} := db_{\max,t} - \Delta_{db,t}$, $lg_{\mathrm{lane},t} := db_{\min,t} + \Delta_{db,t}$. Hence, the final lane constraint is

$$lg_{\mathrm{lane},t} \le a_t^\top p_t \le ug_{\mathrm{lane},t}.$$

3) Chance-Constrained Longitudinal Collision Avoidance: Rear-end safety with the preceding vehicle is enforced in curvilinear coordinates along the common reference path, consistently with the distributed stochastic platooning setting in [17]. Let $\Delta s_{t|k} := s_{t|k}^{\mathrm{front}} - s_{t|k}^{\mathrm{ego}}$, and assume $\Delta s_{t|k} \sim \mathcal{N}(\mu_{\Delta s,t|k}, \sigma_{\Delta s,t|k}^2)$. The safety requirement $\Pr\big(\Delta s_{t|k} \ge d_{\mathrm{safe}}\big) \ge p_{\mathrm{coll}}$ is replaced by the deterministic surrogate

$$\mu_{s,\mathrm{ego},t|k} \le \mu_{s,\mathrm{front},t|k} - \big(d_{\mathrm{safe}} + \beta\sigma_{\Delta s,t|k}\big),\ \beta := \Phi^{-1}(p_{\mathrm{coll}}).$$

Here $\mu_{s,\mathrm{front},t|k}$ is treated as known data obtained from communicated front predictions or assumed trajectories, whereas $\mu_{s,\mathrm{ego},t|k}$ is part of the optimization variables. The ego curvilinear variance is approximated by projecting the planar covariance onto the local tangent direction of the reference path. In the local Frenet geometry, this gives

$$\sigma_{s,\mathrm{ego},t|k}^2 = \frac{\tau^\top \Sigma_{\mathrm{ego},t|k}^{xy}\,\tau}{(1 - \kappa e_c)^2}, \tag{13}$$

where τ is the unit tangent to the reference path, $\kappa = \kappa(s)$ is the local path curvature, and e_c is the contouring error. In implementation, the denominator is numerically clamped to avoid ill-conditioning when $1 - \kappa e_c$ becomes small.

This longitudinal collision logic is an original contribution of the implemented controller.

4) Polyhedral Friction-Circle Constraint: To preserve SQP compatibility, the friction-circle constraint is enforced through the polyhedral inner approximation of [21]. In the implementation, a_x is obtained from the traction model, whereas a_y is approximated as $\kappa(s)v_x^2$; both terms are then linearized at the stage nominal point, so that the resulting friction facets enter the QP as affine constraints.

IV. Simulations and Results

The proposed controller is evaluated under representative operating regimes inspired by the GDU framework [12], without validating runtime supervisor transitions, on a real highway segment of the Great Ring Road (GRA–A90) in Rome, Italy, with a spline-based centerline and $3.6\,\mathrm{m}$-wide lanes. The selected section includes a pronounced curve before a merging area between two road branches. All vehicles share the same road reference and are modeled, in the simulated scenario, as identical passenger cars measuring $4.75 \times 1.8 \times 1.4m$, with mass $m = 1300$kg and wheelbase

$L = 2.8$ m. This homogeneous-vehicle assumption is not critical, since vehicle-specific parameters can be exchanged during platoon formation, while GPR can compensate residual unmodeled dynamics even when a common nominal model is used. Unless otherwise stated, the controller operates with $T_{s,\mathrm{MPC}} = 0.10$ s, prediction horizon $N_p = 20$, and one sparse QP solve per control update.

The leader is treated differently from the followers: its desired progress/velocity profile is assigned exogenously, consistently with the hierarchical platooning architecture in which the global path is generated by upper layers [22], whereas the followers use the cooperative neighbour-based reference law introduced in the previous section. All simulations were run in MATLAB/Simulink 2025b on AMD Ryzen 4500U; each sparse QP required about 0.04 ms on average and less than 0.08 ms in the most critical cases.

A. PLF scenario

The first family of simulations refers to a platooning configuration with leader–predecessor following (PLF) topology that corresponds to the *operational* regime of the GDU logic and constant distance gap policy of $d_0 = 14$ m between vehicle centers of gravity. Since this is the tightest configuration, it represents the most stressful case for the controller, which must guarantee the highest reactivity and the strongest coordination along the string.

1) Nominal cooperative platooning: In this scenario, all vehicles track the common road geometry while preserving the desired inter-vehicle spacing and a coordinated longitudinal evolution along the platoon as shown in Fig. 2, where the leader changes its velocity, causing an oscillation of the follower vehicles.

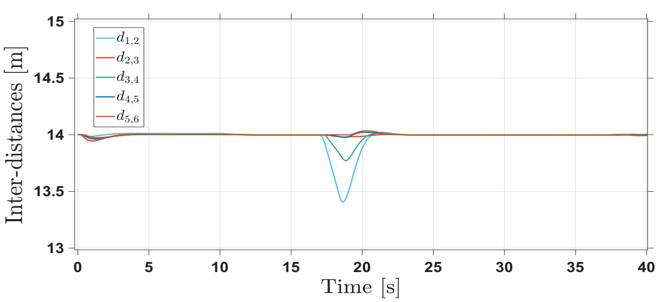

Fig. 2: Inter-vehicle distances. The oscillation is caused by leader-speed change from 22 to 15 m/s.

2) Obstacle-Induced Lane Change: A second platooning test considers a simplified obstacle-induced lane-change manoeuvre. No online high-level planner is implemented; obstacle information and the updated centerline reference are assumed to be provided by the upper-layer architecture. The manoeuvre is thus realized by replacing the reference path, without changing the GPR-SQP-MPCC controller.

Fig. 3 shows the lateral-velocity responses. The lane change starts at about $37\,\mathrm{s}$, while the oscillations observed beforehand are due to the road geometry. The transient caused by the reference update remains bounded and decays after the manoeuvre, confirming stable tracking. The manoeuvre is performed at a commanded speed of $22\,\mathrm{m/s}$ with a constant-distance-gap policy.

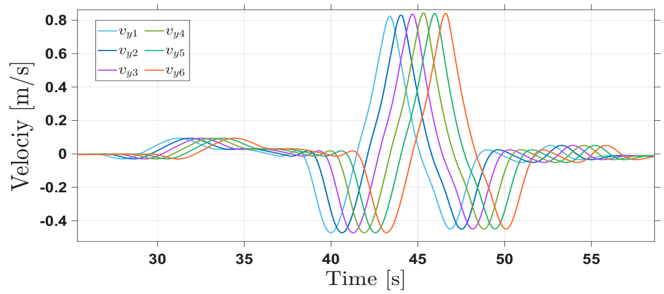

Fig. 3: Lateral-velocity responses during lane change.

3) Emergency braking: The most demanding platooning test is an emergency-braking manoeuvre in which the leader decelerates from 30 m/s to 6 m/s while the platoon travels through a curved road segment.

The contouring penalty is increased with respect to the nominal setting; this prevents the optimizer from recovering feasibility through excessive path deviation during hard braking on a curved segment.

Fig. 4 reports the most representative responses. The longitudinal velocities remain coordinated along the string, the lateral-velocity transients stay bounded, and the inter-vehicle distances remain above the safety threshold. As expected, the strongest transient occurs at braking onset, where path curvature makes the longitudinal–lateral coupling more critical.

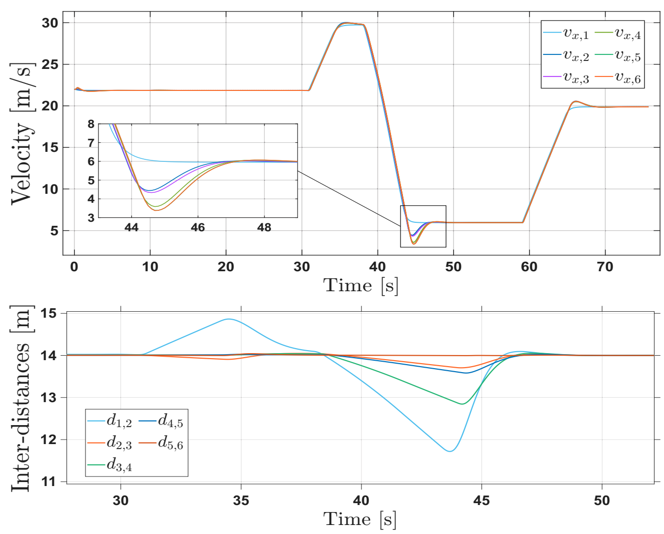

Fig. 4: Longitudinal velocity and inter-vehicle distances.

B. ACC scenario

The second family of simulations refers to the fallback condition corresponding to the *unsafe* regime of the GDU logic. In this case, V2V information can no longer be relied upon as the primary source for predecessor tracking, and the architecture effectively relies on onboard sensing.

In the present implementation, the predecessor state is reconstructed through the simple velocity model proposed in [23], where the prediction is fed by ideal radar-like measurements, while any available V2V information can only be used as an auxiliary correction source. Hence, the cooperative preview available in platooning mode is replaced by a more conservative state-estimation mechanism, consistently with a sensing-based fallback architecture.

Fig. 5 reports the longitudinal-velocity responses in this ACC scenario. The string remains collision-free, but the transient becomes more conservative and a larger overshoot appears toward the rear of the platoon. This behaviour is consistent with the simplified prediction model adopted in the estimator, which reacts less aggressively to abrupt predecessor decelerations than the cooperative preview available in the platooning regime. Therefore, the unsafe case confirms safe fallback operation at the price of reduced performance.

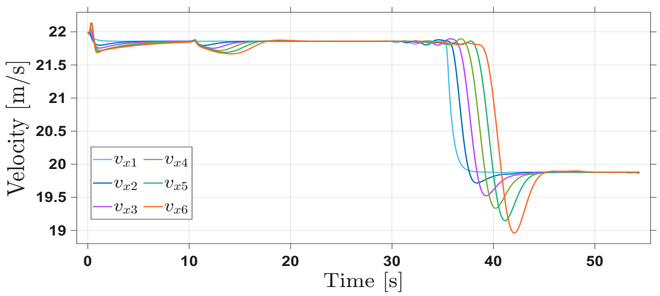

Fig. 5: ACC scenario: longitudinal-velocity response with radar/KF-based predecessor estimation.

Overall, the reported simulations indicate that the proposed GP-SQP-MPCC regulation layer can handle cooperative PLF platooning and sensing-based ACC fallback under the considered operating conditions, while preserving longitudinal safety and satisfactory path-tracking behaviour.

V. Conclusions

This paper presented an implementable distributed communication-aware GPR-SQP-MPCC framework for planar vehicle platooning, combining planar predictive control, GP-based residual compensation, and GDU-inspired outage-aware operation within a unified low-level architecture. The proposed approach explicitly accounts for the coupling between longitudinal and lateral dynamics, while incorporating learning-based model refinement and communication-dependent supervisory logic.

Simulation results in representative multi-vehicle scenarios, including emergency braking, lane changes, and communication loss, demonstrate that the proposed framework can maintain safe inter-vehicle spacing, coordinated motion, and robust path tracking. Moreover, the use of a sparse SQP formulation ensures real-time feasibility, supporting the practical deployment of learning-enhanced predictive control in distributed platooning applications.

Future work will address the integration of covariance-aware SQP schemes for improved uncertainty handling, the development of formally verifiable GDU supervisory mechanisms, and the incorporation of PHY-accurate 5G V2X communication models.

References

[1] K. N. Qureshi and A. H. Abdullah, "A survey on intelligent transportation systems," *Middle-East Journal of Scientific Research*, vol. 15, no. 5, pp. 629–642, 2013.

[2] J. Axelsson, "Safety in vehicle platooning: A systematic literature review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 5, pp. 1033–1045, 2016.

[3] M. Rebelo, S. Rafael, and J. M. Bandeira, "Vehicle platooning: A detailed literature review on environmental impacts and future research directions," *Future Transportation*, vol. 4, no. 2, pp. 591–607, 2024.

[4] H. A. Hamersma and P. S. Els, "Longitudinal vehicle dynamics control for improved vehicle safety," *Journal of Terramechanics*, vol. 54, pp. 19–36, 2014.

[5] S. Liu, X. Chen, Y. Zhao, H. Su, X. Zhou, and K. Zheng, "Comfort-aware lane change planning with exit strategy for autonomous vehicle," *IEEE Transactions on Knowledge and Data Engineering*, vol. 36, no. 7, pp. 2927–2941, 2024.

[6] W. J. Schakel, B. Van Arem, and B. D. Netten, "Effects of cooperative adaptive cruise control on traffic flow stability," in *13th International IEEE Conference on Intelligent Transportation Systems*, pp. 759–764, IEEE, 2010.

[7] X. Liang, W. Gao, C. Hu, and T. Chai, "Cooperative adaptive cruise control of connected and autonomous vehicles via hybrid iteration," *IEEE Transactions on Vehicular Technology*, 2025.

[8] Y. Su, X. Yang, P. Shi, G. Wen, and Z. Xu, "Consensus-based vehicle platoon control under periodic event-triggered strategy," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 54, no. 1, pp. 533–542, 2023.

[9] C. F. Caruntu, C. Braescu, A. Maxim, R. C. Rafaila, and A. Tiganasu, "Distributed model predictive control for vehicle platooning: A brief survey," in *2016 20th International Conference on System Theory, Control and Computing (ICSTCC)*, pp. 644–650, IEEE, 2016.

[10] M. El Zorkany, A. Yasser, and A. I. Galal, "Vehicle to vehicle "v2v" communication: scope, importance, challenges, research directions and future," *The Open Transportation Journal*, vol. 14, no. 1, 2020.

[11] P. S. Narayanan, "Vehicle-to-vehicle (v2v) communication using routing protocols: A review," in *2019 international conference on smart structures and systems (ICSSS)*, pp. 1–10, IEEE, 2019.

[12] S. Hasan, S. Girs, and E. Uhlemann, "Characterization of transient communication outages into states to enable autonomous fault tolerance in vehicle platooning," *IEEE Open Journal of Intelligent Transportation Systems*, vol. 4, pp. 101–129, 2023.

[13] D. Liu, S. Baldi, and S. Hirche, "Collision avoidance in longitudinal platooning: Graceful degradation and adaptive designs," *IEEE Control Systems Letters*, vol. 7, pp. 1694–1699, 2023.

[14] A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1:43 scale rc cars," *Optimal Control Applications and Methods*, vol. 36, p. 628–647, July 2014.

[15] L. Hewing, A. Liniger, and M. N. Zeilinger, "Cautious nmpc with gaussian process dynamics for autonomous miniature race cars," in *2018 European Control Conference (ECC)*, pp. 1341–1348, 2018.

[16] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3363–3370, 2019.

[17] X. Hu, Y. Shi, L. Xie, and H. Su, "Distributed non-linear model predictive control with gaussian process dynamics for two-dimensional motion of vehicle platoon," *IET Intelligent Transport Systems*, vol. 18, no. 1, pp. 1–15, 2024.

[18] MathWorks, "Vehicle Body 3DOF — Vehicle Dynamics Blockset documentation."

[19] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 11 2005.

[20] P. Polcz, T. Péni, and R. Tóth, "Efficient implementation of gaussian process–based predictive control by quadratic programming," *IET Control Theory & Applications*, vol. 17, pp. 968–984, 2023.

[21] R. de Castro et al., "Design of safety-oriented control allocation strategies for overactuated electric vehicles," *Vehicle System Dynamics*, vol. 52, no. 8, pp. 1017–1046, 2014.

[22] P. Varaiya, "Smart cars on smart roads: problems of control," *IEEE Transactions on Automatic Control*, vol. 38, no. 2, pp. 195–207, 1993.

[23] M. Baek, D. Jeong, D. Choi, and S. Lee, "Vehicle trajectory prediction and collision warning via fusion of multisensors and wireless vehicular communications," *Sensors*, vol. 20, no. 1, 2020.