

# Optimization-Based Hierarchical Motion Planning for Autonomous Racing

José L. Vázquez<sup>1,\*</sup>, Marius Brühlmeier<sup>1,\*</sup>, Alexander Liniger<sup>2,\*</sup>, Alisa Rupenyan<sup>1</sup>, John Lygeros<sup>1</sup>

**Abstract**—In this paper we propose a hierarchical controller for autonomous racing where the same vehicle model is used in a two level optimization framework for motion planning. The high-level controller computes a trajectory that minimizes the lap time, and the low-level nonlinear model predictive path following controller tracks the computed trajectory online. Following a computed optimal trajectory avoids online planning and enables fast computational times. The efficiency is further enhanced by the coupling of the two levels through a terminal constraint, computed in the high-level controller. Including this constraint in the real-time optimization level ensures that the prediction horizon can be shortened, while safety is guaranteed. This proves crucial for the experimental validation of the approach on a full size driverless race car. The vehicle in question won two international student racing competitions using the proposed framework; moreover, our hierarchical controller achieved an improvement of 20% in the lap time compared to the state of the art result achieved using a very similar car and track.

## I. INTRODUCTION

Over the past few decades motion planning for autonomous driving has attracted great interest in both academic and industrial research, see for example [1] for a review on motion planning, and [2], [3] for a robotic system perspective. A sub-field of autonomous driving where motion planning is of crucial importance, is autonomous racing, where the goal is to drive around a given race track as fast as possible. This requires a motion planning method that pushes the car to the limit of handling, and is able to handle the nonlinear behavior in this regime [4], [5]. In this paper, we propose a novel hierarchical optimization-based motion planner that exploits the structure of the given task, and demonstrate the performance using the Formula Student Driverless (FSD) car developed at ETH Zurich. The robotic platform with the name *pilatus* is shown in Fig. 1, and won two international autonomous racing competitions using the algorithms presented in this work.

When investigating existing autonomous racing motion planners and controllers, three classes can be distinguished: The first class computes an ideal path and a velocity profile around the track offline, and then follows this path using a static feedback controller [4], [6]. These approaches are intriguing due to their simplicity but since all motion planning is done offline they lack flexibility. The remaining two classes rely on online optimization-based motion planning



Fig. 1. Photo of *pilatus* ©FSG Gosala

[5], [7]. The methods are interesting because they are able to consider the nonlinear vehicle dynamics, track constraints, input constraints, and at the same time plan an optimal motion for the vehicle. The challenge however is that autonomous racing requires fast sampling times in the order of 10-50ms, while finding the optimal trajectory around a race track requires significant foresight in the order of several seconds. This can be challenging since the model complexity and planning horizon both contribute to increase the computational burden. The main difference between the two online optimization based classes lies exactly in addressing this problem. One is based purely on online motion planning. The advantage is that it runs completely online, however, given today's computational platforms, these methods are limited to special applications such as miniature race cars [5], [8]. Extension to full size car is possible [9], but only for limited top speeds. The last approach strikes a balance between the two previous classes and combines both offline and online motion planning [10], [11], [12], [13]. Our work falls under this third class that mixes online and offline elements; the building blocks of the method are outlined below.

In [10] a viability constraint is computed offline that guides the online motion planner. In [13] terminal constraints and cost are estimated using experimental data, these terminal ingredients then guide the online Nonlinear Model Predictive Control (NMPC) motion planner. In [12], a least curvature path as well as velocity profile is computed offline, this trajectory is then tracked with simple NMPC. Finally, [11] computes a least curvature path offline. To follow the path the method first computes a velocity profile that is then used as a terminal velocity constraint in a low-level NMPC motion planner based on a sophisticated vehicle model. Thus, [11] combines the ideas of terminal constraints with an

<sup>1</sup> Automatic Control Lab, ETH Zurich, Switzerland  
{vjose,mariusbr}@student.ethz.ch,  
{ralisa,lygeros}@control.ee.ethz.ch

<sup>2</sup> Computer Vision Lab, ETH Zurich, Switzerland  
alex.liniger@vision.ee.ethz.ch

\* The authors contributed equally.

offline computed ideal line.

Our approach is mostly inspired by [11], specifically, our hierarchical motion planner computes a time-optimal trajectory around the race track offline, and a low-level NMPC follows this path, while restricted by a terminal velocity constraint extracted from the time-optimal trajectory. However, compared to [11] our high-level motion planner computes a full time optimal state-input trajectory using a lap time optimization approach similar to [14], [15]. This makes it possible to use the same model for the two levels, which guarantees that the minimum-time trajectory is drivable while achieving lower lap times. Compared to [11] the model used in the high-level planner is more complex, while that in the lower level is simpler. This resulted in reducing significantly the online computational load, which enables the real-time implementation of our method.

Our work makes the following contributions to the state of the art: First we propose a hierarchical controller where the two levels use the same vehicle model formulated in curvilinear coordinates. Second, we show that for the given FSD setting, following the minimum time trajectory improves performance compared to planning the motion from scratch, and that using terminal velocity constraints allows the use of shorter prediction horizons without sacrificing performance. Finally, we show that the proposed hierarchical controller can race a full-size autonomous FSD race car, finishing fastest in all the attended FSD competitions, and reducing the lap times shown in [9] by 20%, using a similar car and track. See <https://youtu.be/gcnnngFyWnFQ?t=13079> for the run at the formula student Germany competition and <https://youtu.be/vkVBi9LWJo0> for additional experiments and visualizations.

In Section II we introduce the vehicle model and constraints used in our motion planner, and discuss time and space-dependent curvilinear models. In Section III we formulate the high-level lap time optimization problem and the low-level NMPC path following problem. We show numerical and experimental results in Section IV, and conclude the paper and give future direction of research in Section V.

## II. MODEL

We first introduce the vehicle model used in the hierarchical curvilinear motion planner, including the model-related constraints, and the coordinate system. This allows us to highlight similarities and differences in the two layers of our control structure. First the model formulated in curvilinear coordinates is presented, followed by the model-related constraints, and finally the difference between time and space-dependent models is discussed.

The starting point is a bicycle model with dynamic force laws [16]. The main advantage of this model is that it strikes a balance between simplicity and accuracy. It is significantly simpler than a full-fledged double-track model with complex tire forces [11], while still modeling important effects such as tire slip and saturation.

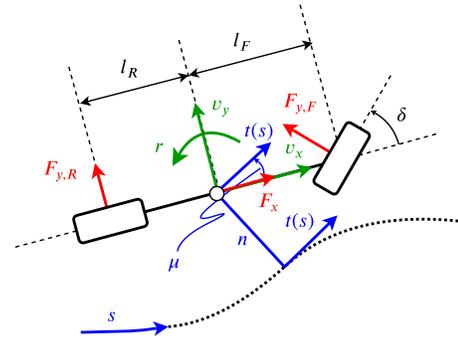


Fig. 2. Illustration of the bicycle model in the curvilinear frame. The figure shows the curvilinear coordinates in blue, the velocity states in green and the forces red. The tangent of the reference path is presented as  $t(s)$ .

### A. Curvilinear Bicycle Model

The bicycle model dynamics is set in curvilinear coordinates, which are formulated locally with respect to a given reference path. In our case the reference path can be the center line of the race track, or an ideal line. By using curvilinear coordinates, no global position or heading is considered, and the state information is given relative to the reference path. The state comprises of the progress (arc-length) along the path  $s$ , the orthogonal deviation from the path  $n$ , and the local heading  $\mu$ . Given the reference path the global coordinates can be recovered, and the curvilinear coordinates can be obtained (locally) by projecting the global coordinates onto the reference path. The bicycle model and the curvilinear coordinates are shown in Fig. 2.

Several assumptions are needed to formulate our vehicle model: (i) the car travels on level ground, (ii) load changes can be neglected<sup>1</sup>, (iii) combined slip can be neglected<sup>2</sup>, (iv) all the drive train forces act at the center of gravity, and (v) the vehicle is a rigid body with mass  $m$  and rotational inertia  $I_z$ . The first assumption is valid as formula student race tracks are level. The remaining assumptions hold since we use a low level traction controller, that distributes the requested drive train force and torque vectoring moment to the four wheels while considering combined slip and load changes, similar to [9].

The dimensions of the vehicle are given by  $l_F$  and  $l_R$  which describe the distance between the CoG and the corresponding front and rear axles. To formulate the dynamics we first consider the curvilinear states  $s$ ,  $n$ , and  $\mu$ , and use the standard approach to integrate the velocities, which relies on the curvature  $\kappa(s)$  of the reference path, see [15] for a detailed discussion. We consider next the longitudinal and lateral velocity  $v_x$  and  $v_y$  and the yaw rate  $r$ , which are influenced by the tire and drive train forces. Finally, we also consider the actuator dynamics by using an integrator model for the steering angle  $\delta$  and the driver command  $T$ . Thus, the state is given by  $\tilde{\mathbf{x}} = [s; n; \mu; v_x; v_y; r; \delta; T]$  and the input as  $\mathbf{u} = [\Delta\delta, \Delta T]$ . Under the assumptions listed above the

<sup>1</sup>Normal loads do not depend on the accelerations of the car.

<sup>2</sup>Lateral and longitudinal forces at the wheels are independent.

dynamics then becomes

$$\begin{aligned}
\dot{s} &= \frac{v_x \cos \mu - v_y \sin \mu}{1 - n\kappa(s)}, \\
\dot{n} &= v_x \sin \mu + v_y \cos \mu, \\
\dot{\mu} &= r - \kappa(s)\dot{s}, \\
\dot{v}_x &= \frac{1}{m}(F_x - F_{y,F} \sin \delta + mv_y r), \\
\dot{v}_y &= \frac{1}{m}(F_{y,R} + F_{y,F} \cos \delta - mv_x r), \\
\dot{r} &= \frac{1}{I_z}(F_{y,F} l_F \cos \delta - F_{y,R} l_R + M_{tw}), \\
\dot{\delta} &= \Delta \delta, \\
\dot{T} &= \Delta T.
\end{aligned} \tag{1}$$

Here  $\kappa(s)$  is the curvature at the progress  $s$ ,  $F_{y,F/R}$  are the lateral tire forces at the front and rear wheel and  $F_x$  is the longitudinal force acting on the car. Finally,  $M_{tw} = p_{tw}(r_t - r)$  is the moment the torque vectoring system produces, where  $p_{tw}$  is the gain of the system and  $r_t = \tan(\delta)v_x/(l_R + l_F)$  is the yaw rate target, see [9] for more details. Note that the curvature contains the full information about the reference path. To simplify the notation we denote the dynamics in (1) by  $\dot{\tilde{\mathbf{x}}} = f_t^c(\tilde{\mathbf{x}}, \mathbf{u})$ , where the superscript  $c$  highlights that it is a continuous time model and the subscript  $t$  that it is a time-domain model.

The tire forces are the most important part of the model, as they model the interaction of the car with the ground. The lateral tire forces  $F_{y,F}$  and  $F_{y,R}$  describe the lateral forces acting on the front, and respectively the rear tires. They are computed using a simplified Pacejka tire model [17],

$$\begin{aligned}
\alpha_F &= \arctan\left(\frac{v_y + l_F r}{v_x}\right) - \delta, \\
\alpha_R &= \arctan\left(\frac{v_y - l_R r}{v_x}\right), \\
F_{y,F} &= F_{N,F} D_F \sin(C_F \arctan(B_F \alpha_F)), \\
F_{y,R} &= F_{N,R} D_R \sin(C_R \arctan(B_R \alpha_R)),
\end{aligned}$$

where  $\alpha_{F/R}$  are the slip angles at the front and, respectively rear wheel, and  $B_{F/R}$ ,  $C_{F/R}$  and  $D_{F/R}$  are the parameters of the simplified Pacejka tire model. Finally, the normal loads on the tires are given by  $F_{N,F} = mgl_R/(l_F + l_R)$  and  $F_{N,R} = mgl_F/(l_F + l_R)$ .

The longitudinal drive train force  $F_x$  is composed of a motor force  $F_M = C_m T$ , which is a linear function of the driver command  $T \in [-1, 1]$ , as well as a constant rolling resistance term  $C_{r0}$  and drag-force term  $C_{r2}v_x^2$ , thus we have  $F_x = C_m T - C_{r0} - C_{r2}v_x^2$ .

## B. Constraints

During the race, it is crucial to ensure that the entirety of the vehicle stays within the track boundaries. This is enforced via heading-dependent constraints on the lateral deviation  $n$ ,

$$\begin{aligned}
n - \frac{L_c}{2} \sin |\mu| + \frac{W_c}{2} \cos \mu &\leq \mathcal{N}_L(s), \\
-n + \frac{L_c}{2} \sin |\mu| + \frac{W_c}{2} \cos \mu &\leq \mathcal{N}_R(s),
\end{aligned} \tag{2}$$

where  $L_c$  is the length and  $W_c$  is the width of the car, and left and right track width at a specific progress  $s$  are given by  $\mathcal{N}_{R/L}(s)$ , similar to [18]. We denote the resulting track constraints (2) as  $\tilde{\mathbf{x}} \in \mathcal{X}_{\text{Track}}$ .

Since the used tire model neglects combined tire forces, we use a friction ellipse constraint to limit the combined tire forces, such that the low level traction controller can better deal with the requested forces. The resulting friction ellipse constraints are given by

$$(\rho_{long} F_{M,F/R})^2 + F_{y,F/R}^2 \leq (\lambda D_{F/R})^2, \tag{3}$$

where  $F_{M,F} = F_{M,R} = F_M/2$  is the motor force at the front, and respectively at the rear wheel, and a static 50/50 force split is assumed between the wheels.  $\rho_{long}$  determines the shape of the ellipse, and controls how much force can be applied in the longitudinal direction, and  $\lambda$  allows to determine the maximum combined force. We denote the friction ellipse constraints in (3) by  $\tilde{\mathbf{x}} \in \mathcal{X}_{\text{FE}}$ .

Finally, we also impose constraints on the physical inputs  $\delta$  and  $T$ , as well as on the rate of change of these inputs  $\Delta \delta$  and  $\Delta T$ . These constraints are simple box constraints and are determined based on physical limits for the steering input and safety considerations for the driving command. For a slim definition we define  $\mathbf{a} = [\delta; T; \Delta \delta; \Delta T]$ , and the box constraints are denoted as  $\mathbf{a} \in \mathcal{A}$ .

## C. Space vs Time Domain

A natural way to formulate dynamical systems is to describe them as a function of time, especially for the purpose of control where control inputs are applied with a fixed sampling frequency. However, for our dynamics it is also possible to formulate them to evolve with respect to space, making the state a function of the progress  $\tilde{\mathbf{x}}(s)$ . This transformation can be performed using the reformulation,

$$\begin{aligned}
\dot{\tilde{\mathbf{x}}} &= \frac{\partial \tilde{\mathbf{x}}}{\partial t} = \frac{\partial \tilde{\mathbf{x}}}{\partial s} \frac{\partial s}{\partial t} \\
&\Rightarrow \frac{\partial \tilde{\mathbf{x}}}{\partial s} = \frac{1}{\dot{s}} f_t^c(\tilde{\mathbf{x}}(s), \mathbf{u}(s)) = f_s^c(\tilde{\mathbf{x}}(s), \mathbf{u}(s))
\end{aligned} \tag{4}$$

Expressing the state as a function of  $s$  comes with several advantages, for example the time is now a function of the states  $t = 1/\dot{s}$ . This allows us to naturally formulate a minimal time problem, as the time it takes to move from  $s = s_0$  to  $s_T$  is given by  $T = \int_{s_0}^{s_T} 1/\dot{s} \, d\sigma$ . One additional advantage is that the curvature does not change based on the inputs, since  $s$  is no longer a function of time. This is especially relevant for long prediction horizons where time dependent formulations need an accurate initial guess of the solution. The last advantage is that the progress state  $s$  becomes redundant and can be dropped from the state vector, thus a reduced state vector can be used for the space-curvilinear model  $\mathbf{x} = [n; \mu; v_x; v_y; r; \delta; T]$ . The associated disadvantages are that first, when discretizing the system it is not time-discrete but space-discrete. This implies that, when using this model for control, inputs should not be changed with a fixed sampling time but at a fixed space interval. Second, pre-multiplying the system with

$1/\dot{s}$  makes the dynamics computationally more complex for an optimization solver, since computing Jacobians becomes more expensive. Finally, the space-domain transformation introduces a singularity if the car stops ( $\dot{s} = 0$ ).

### III. HIERARCHICAL CONTROLLER

One fundamental challenge when implementing an NMPC is the computational burden caused by solving an optimization problem at each time step. The main factors that influence the computation time are the number of states and inputs, and the prediction horizon. By using curvilinear coordinates, the number of states can be reduced compared to the current state of the art method [9]. To reduce the horizon length we first compute the Lap Time Optimization (LTO) trajectory, and then follow this path with a NMPC where the longitudinal speed at the end of the horizon is upper bounded by the LTO speed, see Fig. 3. Note that the LTO path as well as the LTO terminal velocity are fundamental for shortening the horizon.

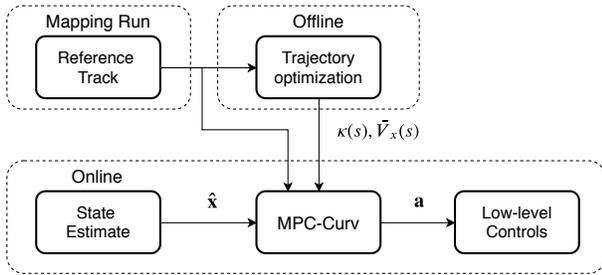


Fig. 3. The hierarchical controller uses the reference track in both stages of the hierarchical controller. First the lap time optimization (LTO) problem computes a reference path, whose curvature  $\kappa(s)$  and speed profile  $\bar{V}_x(s)$  are later used by the NMPC.

An alternative view on this approach is in terms of terminal constraint and recursive feasibility. Since the same model is used for the two levels, and the LTO trajectory is periodic, the LTO trajectory is an invariant set of the system dynamics and the constraints. Thus, if the NMPC exactly steers the system to the LTO trajectory there exists an input sequence that keeps the system within the constraint set indefinitely. The terminal velocity constraint used here is a relaxation, compared to forcing the NMPC to exactly steering the system to the LTO trajectory. Thus, our approach does not guarantee recursive feasibility but comes with some advantages. For one, we avoid a terminal equality constraint that is numerically difficult to deal with. Moreover, the online NMPC has more freedom and can deal with cases where the LTO trajectory is not reachable within the prediction horizon.

To highlight that horizon length is a fundamental issue for autonomous racing, note that [9] required a 2s look ahead at narrow formula student track and a top speed of 15m/s, which is at most a 30m look ahead, and for full-scale cars [11] required a 400m look ahead.

#### A. Lap Time Optimization

The optimal race line is computed by a minimum-time optimal control problem that uses the spatial dynamics (4).

To formulate the optimization problem, we use the center line of the pre-mapped race track as a reference path and discretize the continuous space dynamics. We discretize  $s$  with a discretization distance of  $\Delta s$  and use a spatial forward Euler integrator. Thus, the discrete-space dynamics have the form,

$$\mathbf{x}_{k+1} = f_s^d(\mathbf{x}_k, \mathbf{u}_k) = \mathbf{x}_k + \Delta s f_s^c(\mathbf{x}_k, \mathbf{u}_k).$$

To enforce that the resulting optimal trajectory is periodic, a periodicity constraint is inserted,

$$f_s^d(\mathbf{x}_N, \mathbf{u}_N) = \mathbf{x}_0,$$

where  $N$  is the number of discretization steps, and  $\Delta s$  is chosen such that  $(N+1)\Delta s$  is equal to the arc-length of the center line<sup>3</sup>.

The cost function seeks to minimize the lap time  $T = \sum_{k=0}^N \Delta s / \dot{s}_k$ . The cost also contains two regularization terms, a slip angle cost, which penalizes the difference between the kinematic and dynamic side slip angle, and a penalization of the rate of change of the physical inputs. The first regularization term is

$$B(\mathbf{x}_k) = q_\beta (\beta_{\text{dyn},k} - \beta_{\text{kin},k})^2, \quad (5)$$

where,  $q_\beta$  is a positive weight,  $\beta_{\text{kin},k} = \arctan(\delta_k l_R / (l_F + l_R))$ , and  $\beta_{\text{dyn},k} = \arctan(v_{y,k} / v_{x,k})$ . The regularizer on the rate of change of the physical inputs is  $\mathbf{u}^T R \mathbf{u}$ , where  $R$  is a diagonal weight matrix. In summary, the overall cost function is,

$$j_{\text{LTO}}(\mathbf{x}_k, \mathbf{u}_k) = \Delta s \frac{1}{\dot{s}_k} + \mathbf{u}^T R \mathbf{u} + B(\mathbf{x}_k).$$

Finally, we use the constraints introduced in Section II-B, to formulate the lap time optimization problem,

$$\begin{aligned} \min_{\mathbf{X}, \mathbf{U}} \quad & \sum_{k=0}^N j_{\text{LTO}}(\mathbf{x}_k, \mathbf{u}_k) \\ \text{s.t.} \quad & \mathbf{x}_{k+1} = f_s^d(\mathbf{x}_k, \mathbf{u}_k) \\ & f_s^d(\mathbf{x}_N, \mathbf{u}_N) = \mathbf{x}_0 \\ & \mathbf{x}_k \in \mathcal{X}_{\text{Track}}, \quad \mathbf{x}_k \in \mathcal{X}_{\text{FE}} \\ & \mathbf{a}_k \in \mathcal{A}, \quad k = 0, \dots, N, \end{aligned} \quad (6)$$

where  $\mathbf{X} = [\mathbf{x}_0, \dots, \mathbf{x}_N]$  and  $\mathbf{U} = [\mathbf{u}_0, \dots, \mathbf{u}_N]$ . The problem is formulated in the automatic differentiation package *CppAD* [19] and *Ipopt* [20] is used to solve the resulting nonlinear optimization problem.

#### B. MPC-Curv

We use an online NMPC module that we call MPC-Curv, which follows the LTO path. The formulation of the online MPC-Curv problem is very similar to the offline LTO problem, as they share the goal of finishing a lap as fast as possible. However, MPC-Curv uses the time domain model, since it is better suited for online control. Note that our state estimation pipeline, as well as the low-level controllers

<sup>3</sup>Note that we assume that the time-optimal trajectory does not stop ( $\dot{s} > 0$ ), which is fulfilled for practical problems due to the periodicity constraint.

run in discrete time, additionally, the time domain model is computationally less demanding, as already discussed in II-C. Instead of minimizing time, MPC-Curv maximizes the progress over the horizon  $p = \sum_{t=0}^N \Delta t \dot{s}_t$ , where  $\Delta t$  is the sampling time.

To further reduce the solve times of the MPC-Curv, we aim to reduce the number of states in the optimization problem. Since only the curvature and the track constraints depend on the progress state  $s$ , and  $s$  can be accurately estimated from the previous MPC solution, MPC-Curv fixes  $s$  to the estimated values, and discards it in the dynamics. Note that to get the full benefit from including the  $s$  state, one would need to include a parametric curvature function in the model and constraint evaluation of the MPC which would result in additional computational overhead. Thus, we also use  $\mathbf{x} = [n; \mu; v_x; v_y; r; \delta; T]$  as the state in the MPC-Curv problem.

The MPC-Curv cost function combines the progress optimization with the same regularization terms as in the LTO problem (5). However, since the goal is to follow the LTO path, two path following terms are added, minimizing the deviation and local heading with respect to the reference path. The MPC-Curv cost function can therefore be formulated as

$$j_{\text{MPC}}(\mathbf{x}_t, \mathbf{u}_t) = -\Delta t \dot{s}_t + q_n n_t^2 + q_\mu \mu_t^2 + \mathbf{u}^T R \mathbf{u} + B(\mathbf{x}_t),$$

where,  $q_n$  and  $q_\mu$  are positive path following weights,  $R$  a diagonal matrix with positive weights, and  $B(\mathbf{x})$  is the slip angle cost (5). In addition to the constraints an LTO terminal velocity constraint  $v_{x,T} \leq \bar{V}_x(s_T)$  is added, where  $\bar{V}_x$  comes from the LTO trajectory. Thus the MPC-Curv problem is

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{U}} \quad & \sum_{t=0}^T j_{\text{MPC}}(\mathbf{x}_t, \mathbf{u}_t) \\ \text{s.t.} \quad & \mathbf{x}_0 = \hat{\mathbf{x}} \\ & \mathbf{x}_{t+1} = f_t^d(\mathbf{x}_t, \mathbf{u}_t) \\ & \mathbf{x}_t \in \mathcal{X}_{\text{Track}}, \quad \mathbf{x}_t \in \mathcal{X}_{\text{FE}} \\ & \mathbf{a}_t \in \mathcal{A}, \quad v_{x,T} \leq \bar{V}_x(s_T) \\ & t = 0, \dots, T, \end{aligned}$$

where the subscript  $t$  is used to highlight that the problem is formulated in the time domain. Further,  $\hat{\mathbf{x}}$  is the current curvilinear state estimate and  $T$  is the prediction horizon. The discrete time dynamics  $f_t^d(\mathbf{x}, \mathbf{u})$  is obtained by discretizing the continuous time dynamics with an RK4 integrator. Note that  $s_T$  is fixed and is based on the previous MPC-Curv solution, identical to all other occurrences of the progress state  $s_t$ .

The optimization problem is solved in real-time using ForcesPro, a proprietary interior point method solver optimized for NMPC problems [21], [22].

#### IV. RESULTS

The hierarchical controller was tested in simulation and on-track with the full size autonomous racing vehicle, shown in Fig. 1. The vehicle is the AMZ Driverless 2019 race car, a lightweight single seater race car, that has an all-wheel drive electric powertrain. The race car is equipped with a

complete sensor suite including two LiDARs, two cameras, an optical absolute speed sensor and an INS system, among others. The race car is also equipped with an Intel Xeon E3 that runs the control framework introduced in this paper, but also the mapping, localization and state estimation; see [23], [24], [9] for more details. The following results showcase the performance of the controller, both in simulation and during the final racing event at Formula Student Germany (FSG).

The race track used at the competition is composed of sharp turns, straights and a long sweeping curve. During the competition, the track is first mapped, while a simple pure pursuit controller is used to drive the vehicle. The proposed architecture is then used for all the subsequent runs, based on the map of the track, first the LTO trajectory is computed by solving (6), and online the MPC-Curv is following the LTO path, as shown in Fig. 3

We first describe the implementation details of the LTO problem in Section IV-A, in Section IV-B we show simulation results that investigate the influence of the reference path and the terminal velocity constraint on the performance of the controller. Next, the experimental results using the proposed controller are presented in a competition setting in Section IV-C.

##### A. Computation of the LTO trajectory

For the LTO problem, we use  $N = 1000$  discretization steps, which corresponds to a discretization distance of  $\Delta s = 30.7\text{cm}$ . Finally, we limited the maximum velocity to  $17\text{m/s}$ , for safety reasons. Given this discretization, the optimization problem (6) is solved in approximately 12s on a i7-7500U processor. The resulting LTO trajectory is shown in Fig. 5.

##### B. Following the LTO Trajectory

To validate that an ideal line and terminal velocity constraints are beneficial, we tested three different controllers in simulation. First, we tested MPC-Curv *without* terminal velocity constraint, and using the center line as a reference. To compensate for the lack of terminal constraint, we used a longer look-ahead of 2s, which, given our sampling time of  $\Delta t = 25\text{ms}$  results in  $T = 80$ . Note that this setup is comparable to the MPCC used in [9], however MPC-Curv still comes with computational advantages, due to the reduced number of optimization variables. The second controller follows the LTO path, still without terminal constraint. The required horizon in this case is again of 2s. The current approach proposed here, where MPC-Curv follows the LTO path with a terminal constraint, results in a significantly shorter horizon of 1s. The resulting longitudinal velocity for the three MPC-Curv variants, as well as the LTO velocity profile are shown in Fig. 4. The velocity profiles are shown with respect to the arc-length of the center line, and for reference we highlighted the start-finish line in Fig. 5 which is the  $s = 0$  point.

We can clearly see that following the LTO path allows the MPC-Curv to drive faster, compared to following the center line. We believe that this is due to the fact that following the LTO path partially overcomes the limitations of a finite

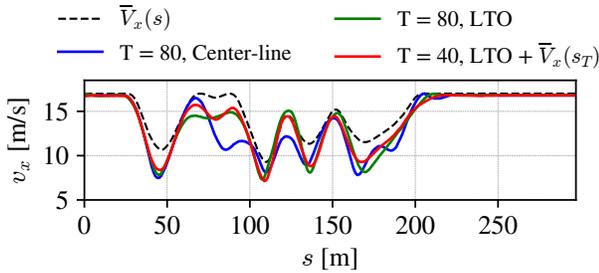


Fig. 4. Comparison of MPC-Curv speed profiles computed for controller following the center line without terminal constraint on  $v_x$  (blue line), controller following the LTO line without terminal constraint (green line), and controller following the LTO line with terminal constraint (red line). The velocity profile computed from LTO is shown for reference (black dashed line).

look ahead, as it offloads line choice to the higher level LTO problem. This allows reaching higher speeds especially in the complex in-field section of the race track. Fig. 4 also clearly shows that using a terminal velocity constraint enables the use of a short 1s look-ahead without sacrificing performance. It comes however with the cost that the brake points are solely dictated by the terminal constraint. The lap times also confirm this discussion, where following the center line results in the slowest lap time of 22.70s, using the LTO path allows for a reduced lap time of 21.67s, with identical lap times for both LTO-based approaches.

Note that from Fig. 4 it seems that the online MPC-Curv is not fully achieving the LTO-based speed. However, the difference is due to different tuning, LTO is tuned to give an optimistic velocity profile and MPC-Curv to drive the real car. Future research is needed to get the online and offline phase of our hierarchical controller closer together.

### C. Experimental Validation

The experimental validation of the hierarchical controller was done during the Formula Student Germany Driverless 2019 competing. The controller at the competition used a look-ahead of 1s with a sampling time of  $\Delta t = 25\text{ms}$ , and a horizon length of  $T = 40$ , identical to the third controller discussed in Section IV-B. The LTO path and velocity profile used as a reference path and terminal velocity constraint are shown in Fig. 5.

One advantage of the proposed control structure is that compared to the current state of the art method [9], the optimization variables per time step are reduced. This allows us to achieve computation times of 12.81ms in average, with only 3% of the time instances over the 25ms limit. Compared to [9] the decrease in computation time is approximately *three-fold*, which is a drastic enhancement, since the same solver and prediction horizon are used. Note that a different processor is used, but even with identical processors, we noticed at least a two-fold decrease in computation times.

Fig. 5 shows that the LTO path (left) and the experimental results (right) are generally very similar. Therefore, MPC-Curv is able to follow the LTO path accurately. The main

difference is that the experimental velocity slightly lags behind the LTO velocity. Since this is not the case in the simulation results, the cause is most likely a mismatch between the model used in the MPC and the dynamics of the real race car. Note that this mismatch also caused the car to slightly miss the apex in the top right curve.

When comparing the time to finish the 10 lap race at the competition, our proposed hierarchical controller achieved an average lap time of 22.63s. This is 20% faster than the average lap time of the 2018 competition winner [9], which achieved a lap time of 28.59s. Note that the track at the two competitions was very similar but not identical and slightly more complex in 2018. Our simulations suggests that the difference between the tracks is in the order of 1.5s. The average lap time was also 8% faster, than the second fastest car at the 2019 competition which achieved an average lap time of 24.49s. We would like to note that 8% in terms of lap time is still a large margin. For a video of the 10 lap race see <https://youtu.be/gcnngFyWnFQ?t=13079>.

Finally, Fig. 6 shows the combined acceleration in lateral and longitudinal direction (GG-diagram) of the run at Formula Student Germany competition, where it can be seen that the car achieves lateral accelerations of up to  $12\text{m/s}^2$ . Compared to the performance of novice race car drivers this can be considered as impressive, however expert drivers would be able to achieve lateral acceleration of over  $18\text{m/s}^2$ .

## V. CONCLUSION

A novel control approach, consisting of a hierarchical MPC using a bicycle model in a curvilinear coordinate system was presented. The high-level controller computes a trajectory that minimizes the lap time, and the low-level nonlinear model predictive path following controller tracks the computed trajectory online. The two levels are further coupled through a terminal constraint, computed in the high-level controller, and used in the online optimization by the low-level controller which ensures that while short prediction horizons can be used online, safety is maintained at all times. The framework was tested on a full-size Formula Student race car, demonstrating significant performance improvements in computation time and lap time reduction, as compared to the current state of the art results achieved on similar platform and track. Further research will address improvements in the coupling of the higher and the lower level controllers, and investigating model learning [25] which we believe will bring us close to human expert performance.

## ACKNOWLEDGEMENT

We would like to thank the entire AMZ Driverless team, this work would not have been possible without the effort of every single member, we are glad for having the opportunity to work among such amazing people.

## REFERENCES

- [1] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *Trans. Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, 2016.

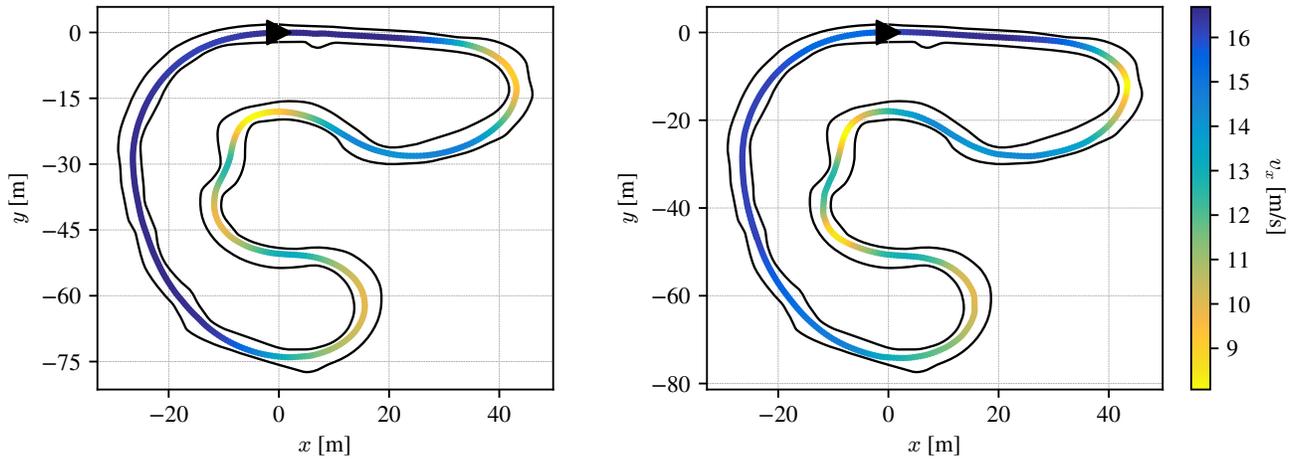


Fig. 5. The racing trajectory calculated by the lap time optimization (left panel) and the real-time trajectory of the vehicle (right panel). The information about  $v_x$  is given by the color of the trajectory.

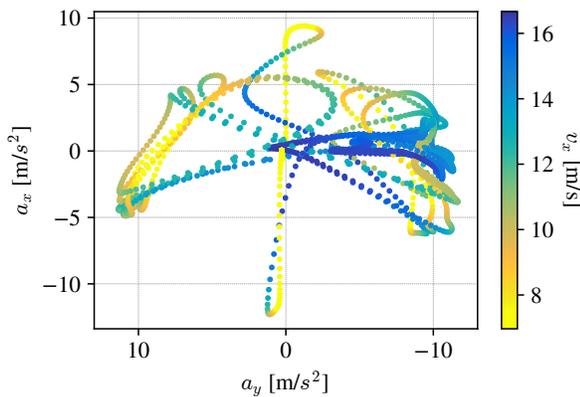


Fig. 6. Longitudinal and lateral accelerations during Formula Student Germany.

[2] M. Buehler, K. Iagnemma, and S. Singh, *The 2005 DARPA grand challenge: the great robot race*. Springer, 2007.

[3] —, *The DARPA urban challenge: autonomous vehicles in city traffic*. Springer, 2009.

[4] K. Kritayakirana and J. C. Gerdes, “Autonomous vehicle control at the limits of handling,” *International Journal of Vehicle Autonomous Systems*, vol. 10, no. 4, pp. 271–296, 2012.

[5] A. Liniger, A. Domahidi, and M. Morari, “Optimization-based autonomous racing of 1:43 scale RC cars,” *Optimal Control Applications and Methods*, vol. 36, no. 5, pp. 628–647, 2015.

[6] J. Betz, A. Wischnewski, A. Heilmeier, F. Nobis, L. Hermansdorfer, T. Stahl, T. Herrmann, and M. Lienkamp, “A software architecture for the dynamic path planning of an autonomous racecar at the limits of handling,” in *International Conference on Connected Vehicles and Expo (ICCVE)*, 2019.

[7] J. Funke, M. Brown, S. M. Erlen, and J. C. Gerdes, “Collision avoidance and stabilization for autonomous vehicles in emergency scenarios,” *IEEE Transactions on Control Systems Technology*, vol. 25, no. 4, pp. 1204–1216, 2016.

[8] R. Verschuere, M. Zanon, R. Quirynen, and M. Diehl, “Time-optimal race car driving using an online exact hessian based nonlinear mpc algorithm,” in *2016 European Control Conference (ECC)*, 2016.

[9] J. Kabzan *et al.*, “AMZ Driverless: The full autonomous racing system,” *arXiv:1905.05150*, 2019.

[10] A. Liniger and J. Lygeros, “Real-time control for autonomous racing

based on viability theory,” *IEEE Transactions on Control Systems Technology*, vol. 27, no. 2, pp. 464–478, 2019.

[11] T. Novi, A. Liniger, R. Capitani, and C. Annicchiarico, “Real-time control for at-limit handling driving on a predefined path,” *Vehicle System Dynamics*, 2019.

[12] D. Caporale, A. Settimi, F. Massa, F. Amerotti, A. Corti, A. Fagiolini, M. Guiggian, A. Bicchi, and L. Pallottino, “Towards the design of robotic drivers for full-scale self-driving racing cars,” in *International Conference on Robotics and Automation (ICRA)*, 2019.

[13] U. Rosolia, A. Carvalho, and F. Borrelli, “Autonomous racing using learning model predictive control,” in *American Control Conference (ACC)*, 2017.

[14] R. Lot and F. Biral, “A curvilinear abscissa approach for the lap time optimization of racing vehicles,” *IFAC World Congress*, 2014.

[15] A. Rucco, G. Notarstefano, and J. Hauser, “An efficient minimum-time trajectory generation strategy for two-track car vehicles,” *IEEE Transactions on Control Systems Technology*, vol. 23, no. 4, pp. 1505–1519, July 2015.

[16] R. N. Jazar, *Vehicle Dynamics: Theory and Application*. Springer, 2008.

[17] H. B. Pacejka and E. Bakker, “The magic formula tyre model,” *Vehicle system dynamics*, vol. 21, no. S1, pp. 1–18, 1992.

[18] A. Liniger and L. Van Gool, “Safe motion planning for autonomous driving using an adversarial road model,” *arXiv preprint arXiv:2005.07691*, 2020.

[19] B. Bell. (2019) Cppad: A package for c++ algorithmic differentiation. [Online]. Available: <http://www.coin-or.org/CppAD>

[20] A. Wächter and L. T. Biegler, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.

[21] A. Domahidi and J. Jerez, “Forces professional,” Embotech AG, url=<https://embotech.com/FORCES-Pro>, 2014–2019.

[22] A. Zanelli, A. Domahidi, J. Jerez, and M. Morari, “Forces nlp: an efficient implementation of interior-point methods for multistage nonlinear nonconvex programs,” *International Journal of Control*, vol. 93, no. 1, pp. 13–29, 2020.

[23] M. I. Valls, H. F. Hendriks, V. J. Reijgwart, F. V. Meier, I. Sa, R. Dubé, A. Gawel, M. Bürki, and R. Siegart, “Design of an autonomous racecar: Perception, state estimation and system integration,” in *International Conference on Robotics and Automation (ICRA)*, 2018.

[24] N. Gosala, A. Bühler, M. Prajapat, C. Ehmke, M. Gupta, R. Sivanesan, A. Gawel, M. Pfeiffer, M. Bürki, I. Sa *et al.*, “Redundant perception and state estimation for reliable autonomous racing,” in *International Conference on Robotics and Automation (ICRA)*, 2019.

[25] 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.