Parallel autonomy in automated vehicles
Safe motion generation with minimal intervention

Schwarting, Wilko; Alonso-Mora, Javier; Pauli, Liam; Karaman, Sertac; Rus, Daniela

DOI

[10.1109/ICRA.2017.7989224](https://doi.org/10.1109/ICRA.2017.7989224)

Publication date

2017

Document Version

Accepted author manuscript

Published in

Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2017)

Citation (APA)

Schwarting, W., Alonso-Mora, J., Pauli, L., Karaman, S., & Rus, D. (2017). Parallel autonomy in automated vehicles: Safe motion generation with minimal intervention. In I-M. Chen, & Y. Nakamura (Eds.), *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2017)* (pp. 1928-1935). IEEE . <https://doi.org/10.1109/ICRA.2017.7989224>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Parallel Autonomy in Automated Vehicles: Safe Motion Generation with Minimal Intervention

Wilko Schwarting¹, Javier Alonso-Mora^{1,2}, Liam Paull¹, Sertac Karaman³, Daniela Rus¹

Abstract—Current state-of-the-art vehicle safety systems, such as assistive braking or automatic lane following, are still only able to help in relatively simple driving situations. We introduce a Parallel Autonomy shared-control framework that produces safe trajectories based on human inputs even in much more complex driving scenarios, such as those commonly encountered in an urban setting. We minimize the deviation from the human inputs while ensuring safety via a set of collision avoidance constraints. We develop a receding horizon planner formulated as a Non-linear Model Predictive Control (NMPC) including analytic descriptions of road boundaries, and the configurations and future uncertainties of other traffic participants, and directly supplying them to the optimizer without linearization. The NMPC operates over *both* steering and acceleration simultaneously. Furthermore, the proposed receding horizon planner also applies to fully autonomous vehicles. We validate the proposed approach through simulations in a wide variety of complex driving scenarios such as left-turns across traffic, passing on busy streets, and under dynamic constraints in sharp turns on a race track.

I. INTRODUCTION

Globally, over 3000 people are killed *every day* [1] in vehicle-related accidents and over one hundred thousand are injured or disabled on average. Worse still is that this number is continuing to increase [2]. In the United States, 11% of accidents are caused by driver distraction (such as cell phone use), 31% involve an impaired driver due to alcohol consumption, 28% involved speeding, and an additional 2.6% were due to fatigue [3]. This troubling trend has resulted in the continued development of advanced safety systems by commercial car manufacturers. For example, systems exist to automatically brake in the case of unexpected obstacles [4], maintain a car in a lane at a given speed, alert users of pedestrians, signage, and other vehicles on the roadway [5]. However, the scenarios that these systems are able to deal with are relatively simple compared to the diverse and complicated situations that we find ourselves in as human drivers routinely. In this work we propose a framework for advanced safety in complex scenarios that we refer to as Parallel Autonomy, which minimizes the deviation from the human input while ensuring safety. The design of the system has two main objectives: (a) minimal intervention - we only apply autonomous control when necessary, and (b) guaranteed safety - the collision free state of the vehicle is explicitly enforced through constraints in the optimization. Although the focus is on cars on roads one can easily apply the method to other domains in robotics.

¹Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT, Cambridge, MA, USA {wilkos, jalonsom, lpaull, rus}@csail.mit.edu

²Delft Center for Systems and Control, Delft University of Technology, Delft, Netherlands j.alonsomora@tudelft.nl

³Laboratory of Information and Decision Systems (LIDS), MIT, Cambridge, MA, USA sertac@mit.edu

Toyota Research Institute ("TRI") provided funds to assist the authors with their research but this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity.

Fig. 1. Parallel Autonomy in complex driving scenarios: Human driver (red) tries to accelerate into an intersection, as shown by the red bar in the lower left inset. However, given the future uncertainty of the other vehicles positions the Parallel Autonomy system prevents the vehicle from continuing and potentially inhibits a collision.

We provide a formulation and algorithmic solution to Parallel Autonomy based on a non-linear Model Predictive Control (NMPC) policy, under the assumptions of known current location of the ego vehicle, the road boundaries, and of other vehicles on the road. Uncertain predictions of future vehicle states in the form of a posterior distribution are parameterized by their means and covariances, which are assumed to be available, e.g. from an inference framework. Specifically, we:

- Incorporate the time-varying uncertainty related to the dynamic obstacle predictions explicitly in the optimization,
- Follow the road by contour tracking within the shared control paradigm, introduce additional constraints for the road boundaries and dynamic obstacles, while maintaining the ability to plan over long time horizons ($\sim 9s$),
- Employ a non-linear model of the vehicle and simultaneously optimize over steering and acceleration.

The basic operation of the controller is shown in Fig. 1, where the driver attempts to cut in front of oncoming traffic to make a left turn, however the Parallel Autonomy system prevents this action to avoid a collision with the oncoming vehicles.

This paper contributes:

- A formulation of Parallel Autonomy as a shared control approach between humans and robots that adheres to the minimal intervention principle and is able to handle complex driving scenarios.
- The development of a real-time NMPC that operates over both speed and steering, and long time-horizons.

In Sec. II we summarize the related work in the field whereas in Sec. III we present the Parallel Autonomy control approach. In Sec. IV we provide a concrete instantiation of the framework and present the NMPC approach to solve it. Finally, we show detailed simulation results and conclusions in Sec. V and Sec. VI respectively.

II. RELATED WORKS

In this section we will provide an overview of the related work in the areas of shared control for autonomous vehicles and Model

Predictive Control (MPC).

A. Shared Control of Autonomous Vehicles

In theory, safety can be guaranteed if we can compute the set of the states for which the vehicle will *inevitably* have a collision and then ensure that we never enter that set. The set is referred to by different terms in the literature, such as the capture set [6], [7], the inevitable collision states (ICS) [8], [9], [10], and the target set [11]. However, without some assumptions or limiting the applicability to relatively simplistic scenarios, this set is difficult to compute analytically. These ICS inspired methods tend to (a) only intervene when the system is at the boundary of the capture set, which can cause undesirable behavior and (b) toggle between either the autonomous system input or the human input. We will follow the idea of [6], [9], [10] and will define a set of probabilistic constraints for collision avoidance.

In this work we directly incorporate the human inputs into an optimization framework in a minimally invasive manner and also add a soft nudging behavior to guide the driver. One of our key objectives is to minimize the amount of deviation of the autonomous system's plan from the driver's intent. This minimization approach has also been formulated for driving applications in various ways in the literature: Shia et al. [12] directly minimize the difference of the steering angle necessary to achieve safe trajectories and the human predicted input, and, similarly, Gao et al. [13] minimize the difference in steering wheel angle only. Erlien et al. [14] minimize the deviation from desired front wheel lateral force with an additional discount factor with increasing time. In contrast, our approach is capable of controlling steering velocity and acceleration simultaneously. Alonso-Mora et al. [15] minimize the deviation from human inputs, in this case orientation and speed, via a convex constrained optimization to generate safe motion of a wheelchair using velocity obstacles. We minimize the (weighted) difference between the human and autonomous system's control inputs jointly in both steering and acceleration, and are able to blend in additional trajectory-specific costs, while strictly enforcing the safety constraints.

B. Model Predictive Control

A variety of MPC approaches applied to shared control for vehicles exist in the literature. For example, Gray et al. [16] use a hierarchical MPC approach for motion primitive based path planning and path tracking that switches control to and from the driver as a function of driver attentiveness to avoid static obstacles. Anderson et al. [17] employ a constrained pathless MPC approach blending human and controller inputs based on a trajectory-criticality function controlling steering commands only. Erlien et al. [14] define vehicle-stability and environmental envelopes to supply safe steering commands at constant speed in a discretized environment. Gao et al. [13] use robust NMPC to avoid only static obstacles while tracking the roads center line over a very short horizon of less than 1.5s. In contrast, our approach can handle complex road scenarios with dynamic maneuvers and obstacles, and to some extent uncertain environments with steering and acceleration control over long horizons

For most related MPC methods in the literature [18], [17], [19], [14] time dependent cost functions, and road constraints need to be specified pre-optimization for specific time steps, or a fixed path is generated and tracked [13]. The resulting divergence from the initial conditions of the optimization can yield invalid linearized constraints and unpredictable planning behavior. We exploit recent advances in efficient Interior-Point solvers [20] and

directly solve the NMPC problem instead, focusing on making all costs and constraints available to the solver without manual linearization.

Model Predictive Contouring Control (MPCC) [21], [22] relaxes the timing and path constraint by parametrizing costs and constraints by path progress instead of time inside a corridor. The formulation is analogously applicable to vehicles following roads [23].

III. PROBLEM FORMULATION

The Parallel Autonomy problem is based on two overarching principles.

- *Minimal intervention* with respect to the human driver: the control inputs to the vehicle should be as close as possible to those of the human driver.
- *Safety*: The probability of collision with respect to the environment and other traffic participants is below a given threshold.

A. Definitions

We use the discrete time shorthand $k \triangleq t_k$, where $t_k = t_0 + \sum_{i=1}^k \Delta t_i$, with t_0 the current time and Δt_i the i -th timestep of the planner. Vectors are bold.

1) *Ego vehicle*: At time k , we denote the configuration of the ego-vehicle, typically position $\mathbf{p}_k = (x_k, y_k)$, linear velocity v_k , orientation ϕ_k and steering angle δ_k , by $\mathbf{z}_k = [\mathbf{p}_k, \phi_k, \delta_k, v_k] \in \mathcal{Z}$. Its control input, typically steering velocity $\dot{\delta}_k$ and acceleration a_k , is labeled $\mathbf{u}_k = [u_k^\delta, u_k^a] \in \mathcal{U}$.

The evolution of the state of a vehicle is then represented by a general discrete dynamical system

$$\mathbf{z}_{k+1} = f(\mathbf{z}_k, \mathbf{u}_k), \quad (1)$$

described in Sec. IV-B.

Let $\mathcal{B}(\mathbf{z}_k) \subset \mathbb{R}^2$ be the area occupied by the ego-vehicle at state $\mathbf{z}_k$. In particular, we model it as a union of circles as shown in Fig. 3.

2) *Other traffic participants*: Other traffic participants, such as vehicles, pedestrians and bikes, are indexed by $i = \{1, \dots, n\}$. Their configuration and control input are given by $\mathbf{z}_k^i \in \mathcal{Z}_i$ and $\mathbf{u}_k^i \in \mathcal{U}_i$. To incorporate uncertainty, we assume a posterior distribution that describes the future state of the vehicles for up to m timesteps is available, e.g. from an inference framework. The distributions are parametrized by their mean state $\mathbf{z}_{1:m}^i$ and covariance $\sigma_{1:m}^i$. High uncertainty in prediction can therefore be reflected in the covariance $\sigma_{1:m}^i$.

At a given state, each traffic participant occupies an area $\mathcal{B}^i(\mathbf{z}_k^i, \sigma_k^i, p_\epsilon) \subset \mathbb{R}^2$ with probability larger than p_ϵ . Here p_ϵ is the accepted probability of collision. We model them as ellipses that grow in size with uncertainty, as described in the forthcoming Sec. IV-E.

3) *Free space*: We consider the workspace $\mathcal{W} = \mathbb{R}^2$ and an obstacle map $\mathcal{O} \subset \mathcal{W}$ containing the static obstacles, such as the limits of the road. We define the environment $\mathcal{E}(k)$ as the state of the world (obstacles, traffic participants) at a time instance k .

B. Parallel Autonomy

We formulate a general discrete time constrained optimization with m timesteps, with time horizon $\tau = \sum_{k=1}^m \Delta t_k$. We use the following notation for a set of states $\mathbf{z}_{0:m} = [\mathbf{z}_0, \dots, \mathbf{z}_m] \in \mathcal{Z}^{m+1}$ and for a set of inputs $\mathbf{u}_{0:m-1} = [\mathbf{u}_0, \dots, \mathbf{u}_{m-1}] \in \mathcal{U}^m$.

The objective is to compute the optimal inputs $\mathbf{u}_{0:m-1}^*$ for the ego-vehicle that minimize a cost function $\hat{J}_h(\mathbf{u}_{0:m-1}, \mathbf{u}_0^h) + \hat{J}_t(\mathbf{z}_{0:m}, \mathbf{u}_{0:m-1})$, where

- $\hat{J}_h(\mathbf{u}_{0:m-1}, \mathbf{u}_0^h)$ is a cost term that minimizes the deviation from the currently observable human input $\mathbf{u}_0^h$.
- $\hat{J}_t(\mathbf{z}_{0:m}, \mathbf{u}_{0:m-1})$ is a cost term that only depends on intrinsic properties of the planned trajectory. It can include various optimization objectives such as energy minimization, comfort, or following a lane.

The optimization is subject to a set of constraints that represent: (1) the transition model of the ego vehicle, (2) no collisions with the static obstacles and (3) no collisions with other traffic participants up to probability p_ϵ .

Given the posterior, parametrized by $\mathbf{z}_{0:m}^i$ and $\sigma_{1:m}^i$, for all traffic participants $i = 1, \dots, n$ and the initial state $\mathbf{z}_0$ of the ego vehicle, the optimal trajectory for the ego vehicle is then given by

$$\begin{aligned} \mathbf{u}_{0:m-1}^* &= \arg \min_{\mathbf{u}_{0:m-1}} \hat{J}_h(\mathbf{u}_{0:m-1}, \mathbf{u}_0^h) + \hat{J}_t(\mathbf{z}_{0:m}, \mathbf{u}_{0:m-1}) \\ \text{s.t. } \mathbf{z}_{k+1} &= f(\mathbf{z}_k, \mathbf{u}_k) \\ \mathcal{B}(\mathbf{z}_k) \cap \mathcal{O} &= \emptyset \\ \mathcal{B}(\mathbf{z}_k) \cap \bigcup_{i \in \{1, \dots, n\}} \mathcal{B}^i(\mathbf{z}_k^i, \sigma_k^i, p_\epsilon) &= \emptyset \\ &\quad \forall k \in \{0, \dots, m\}. \end{aligned} \quad (2)$$

IV. METHOD

In this section we describe the method to solve Eq. (2).

A. Overview

We formulate a NMPC to compute a safe trajectory for the predefined time horizon. The constrained optimization consists of the following costs and constraints.

1) *Cost*: To maintain generality of the problem formulation while easing the understanding of the specifics of the instantiation, the notation of $\hat{J}_h$, $\hat{J}_t$ and Eq. (25) will be slightly altered to J_h , J_t , cf. Eq. (24).

The cost term J_h is given by the deviation from the acceleration and steering angle specified by the human driver. This term is described in Sec. IV-F.

The cost term J_t is defined in Sec. IV-G and consists of terms responsible for giving feedback to the driver if diverted too far from the road's center in the form of slightly *nudging* the driver back into the correct direction without strong intervention. Another term encodes making progress along a reference path—typically the middle of the current lane, and one to improve smoothness of the trajectory.

2) *Constraints*: The optimization is subject to a set of constraints: (1) to respect the transition model of the system, described in Sec. IV-B, (2) to maintain the vehicle within the limits of the road, indicated in Sec. IV-D and (3) to avoid other traffic participants in the sense of guaranteeing a probability of collision below p_ϵ , as given in Sec. IV-E.

3) *Constrained Optimization*: Since we do not currently have a prediction over future driver commands, we propose a linear combination between the cost J_h for minimal intervention and the trajectory cost J_t . At the planning time, full weight is given to the minimal intervention cost J_h and close to zero weight to J_t . As planning time progresses the impact of J_t increases and the weight of J_h decreases. The resulting MPC, which solves Eq. (2), including the specific combination of costs is described in Sec. IV-H.

B. Motion Model

Previous approaches utilized constant longitudinal speed and small angle assumptions [14], [19], [17] in selected static obstacle avoidance scenarios along simple straight roads. In contrast we will consider the impact of longitudinal speed control for higher safety in dynamic, more general and more complex traffic environments.

The MPC's motion model is a simplified car model with a fixed rear wheel and a steerable front wheel with state $\mathbf{z}$ and controls $\mathbf{u}$ as defined in Sec. III-A.1. The rear-wheel driven vehicle with inter-axle distance L and continuous kinematic model

$$\underbrace{\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\phi} \\ \dot{\delta} \\ \dot{v} \end{bmatrix}}_{\dot{\mathbf{z}}} = \begin{bmatrix} v \cos(\phi) \\ v \sin(\phi) \\ \frac{v}{L} \tan(\delta) \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \underbrace{\begin{bmatrix} u^\delta \\ u^a \end{bmatrix}}_{\mathbf{u}}, \quad (3)$$

is described by a discrete time model by integration $\mathbf{z}_{k+1} = \mathbf{z}_k + \int_k^{k+\Delta t_k} \dot{\mathbf{z}} dt = f(\mathbf{z}_k, \mathbf{u}_k)$.

We limit steering angle, $\|\delta\| \leq \delta_{\max}$, steering speed, $\|u^\delta\| \leq \dot{\delta}_{\max}$, and longitudinal speed, $v \leq v_{\max}$, to reasonable values conforming to vehicle performance and the rules of the road, e.g. speed limits.

We will account for and prohibit unsafe driving modes such as high speeds in sharp turns by limiting the yaw-rate $\|\dot{\phi}\| \leq \dot{\phi}_{\max}$, as well as extreme breaking and accelerations $a_{\min} \leq u^a \leq a_{\max}$. As a result, slip is assumed to be sufficiently limited due to reasonably less aggressive driving behavior. The modification is in line with our main goal: driver safety.

C. Nonlinear Model Predictive Contouring Control

In this section we build on the MPCC method of [21], [22], [23] and apply it to our problem setting. The MPCC approach is a good choice for our parallel autonomy formulation since we don't need to enforce that the vehicle *exactly* follows a reference trajectory or path, but instead stays within the corridor of safety limits.

1) *Progress on Reference Path*: The vehicle at position (x_k, y_k) at time k tracks a continuously differentiable and bounded two-dimensional geometric reference path $(x^P(\theta), y^P(\theta))$ of path parameter θ with

$$\mathbf{t} = \begin{bmatrix} \frac{\partial x^P(\theta)}{\partial \theta} \\ \frac{\partial y^P(\theta)}{\partial \theta} \end{bmatrix}, \quad \mathbf{n} = \begin{bmatrix} -\frac{\partial y^P(\theta)}{\partial \theta} \\ \frac{\partial x^P(\theta)}{\partial \theta} \end{bmatrix} \quad (4)$$

being the tangential and normal vectors.

The heading of the path is described by:

$$\phi^P(\theta_k) = \arctan \left(\frac{\partial y^P(\theta)}{\partial x^P(\theta)} \right). \quad (5)$$

The path is parametrized by the arc-length ($\partial\theta/\partial s = 1$) allowing us to estimate the progress of the vehicle with velocity v_k along the reference path along the vehicle's actual path $s = \int v dt$. While parametrization of curves by the arc-length is not trivial, if the distance between knots is small in relation to their curvature, spline parametrization is close to the arc-length. Since our vehicle will follow a given road with sufficiently low deviation from the reference, enforced by the road's boundary, we can assume that

$$\Delta\theta \approx \Delta s = v\Delta t \quad (6)$$

holds. This additional assumption yields an approximated progress along the path parameter

$$\theta_{k+1} = \theta_k + v_k \Delta t_k \quad (7)$$

where $v_k \Delta t_k$ describes the approximated progress at time step k . Ideally, we want to compute the path parameter $\theta^P(x_k, y_k)$ of the closest point on the reference path to (x_k, y_k) . Finding $\theta^P(x_k, y_k)$ analytically is infeasible in the general case, which makes the direct projection operator unsuitable for fast optimization. Therefore, $\theta^P(x_k, y_k)$ is approximated by Eq. (7).

Fig. 2. Approximation of actual path abscissa θ^P by virtual integrator θ_k , and resulting approximation of lag error e_k^l by $\tilde{e}^l(\mathbf{z}_k, \theta_k)$, and contouring error e_k^c by $\tilde{e}^c(\mathbf{z}_k, \theta_k)$. $\tilde{e}^c(\mathbf{z}_k, \theta_k)$ is also used for an approximation of the lateral distance of the vehicle to the reference path.

2) *Longitudinal Error*: The approximation of the curvilinear abscissa $\theta^P(x_k, y_k)$ by θ_k introduces errors (cf. Fig. 2), if the vehicle's actual path deviates from the reference path. A first order approximation of the resulting error of this approximation in longitudinal direction of the vehicle with respect to the reference path yields the lag error

$$\tilde{e}^l(\mathbf{z}_k, \theta_k) = \frac{\mathbf{t}_k^T}{\|\mathbf{t}_k\|} \begin{bmatrix} x_k - x^P(\theta_k) \\ y_k - y^P(\theta_k) \end{bmatrix} \quad (8)$$

$$= -\cos \phi^P(\theta_k) (x_k - x^P(\theta_k)) \quad (9)$$

$$- \sin \phi^P(\theta_k) (y_k - y^P(\theta_k)) \quad (10)$$

projecting the position error of the vehicle with respect to the path's abscissa θ_k along the path's tangent $\mathbf{t}_k$, see Fig. 3.

For sufficiently small $\tilde{e}^l(\mathbf{z}_k, \theta_k)$ the approximated path progress is close to the actual curvilinear abscissa (Eq. (6)), and $\theta_k \approx \theta^P(x_k, y_k)$. The lag error $\tilde{e}^l(\mathbf{z}_k, \theta_k)$ needs to be strongly penalized in the MPCC optimization to keep the error of the approximated evolution θ_k along the path sufficiently small.

3) *Contouring Error*: The deviation of the vehicle's actual position from the estimated position is projected onto the path's normal and is called contouring error:

$$\tilde{e}^c(\mathbf{z}_k, \theta_k) = \frac{\mathbf{n}_k^T}{\|\mathbf{n}_k\|} \begin{bmatrix} x_k - x^P(\theta_k) \\ y_k - y^P(\theta_k) \end{bmatrix} \quad (11)$$

$$= \sin \phi^P(\theta_k) (x_k - x^P(\theta_k)) \quad (12)$$

$$- \cos \phi^P(\theta_k) (y_k - y^P(\theta_k)) \quad (13)$$

It is thus a good measure of how far the vehicle deviates from a given reference path.

The MPCC cost function

$$J_{\text{MPCC}}(\mathbf{z}_k, \theta_k) = \mathbf{e}_k^T \mathbf{Q} \mathbf{e}_k - v_k \quad (14)$$

with path error vector formed from lag and contouring error

$$\mathbf{e}_k = \begin{bmatrix} \tilde{e}^l(\mathbf{z}_k, \theta_k) \\ \tilde{e}^c(\mathbf{z}_k, \theta_k) \end{bmatrix} \quad (15)$$

balances the trade-off between contouring error, lag error, and approximated path progress v_k .

D. Road Representation

All vehicles' reference paths are parametrized by C^1 -continuous clothoids following the road network through pre-specified points. We approximate the clothoids by cubic-splines of closely spaced knots parametrizing the spline by the arc-length to sufficient accuracy. In contrast to computationally expensive evaluation of clothoids, cubic splines provide an analytical parametrization of the reference path, boundaries of the road, and their derivatives needed for solving the nonlinear optimization.

The signed lateral distance $d(\mathbf{z}_k, \theta)$ of the vehicle's position (x_k, y_k) to the reference path is given by the projection along the normal of the reference path at the actual curvilinear abscissa θ^P , again approximated by θ_k such that $d(\mathbf{z}_k, \theta_k) = \tilde{e}^c(\mathbf{z}_k, \theta_k)$.

The free and drivable space of the ego vehicle at the path abscissa θ_k is limited by both the left road boundary $b_l(\theta_k)$ and the right road boundary $b_r(\theta_k)$ which are parametrized by cubic splines to enable analytic evaluation and derivation. The boundaries may also enclose other static obstacles.

The ego vehicle's lateral offset to the path is limited by

$$b_l(\theta_k) + w_{\max} \leq d(\mathbf{z}_k, \theta_k) \leq b_r(\theta_k) - w_{\max} \quad (16)$$

where $w_{\max}$ is an upper bound on the vehicle's outline projected onto the reference path's normal. $w_{\max}$ is larger than half the vehicle's width, since the ego vehicle's relative orientation to the path needs to be accounted for, e.g. when it turns. We constrain the difference between the ego vehicle's heading ϕ_k and the path's heading $\phi^P(\theta_k)$

$$\|\phi_k - \phi^P(\theta_k)\| \leq \Delta \phi_{\max} \quad (17)$$

to maintain validity of $w_{\max}$ as an upper bound. Simply taking the vehicle's radius as an upper bound turned out to be too conservative.

Fig. 3. Ego vehicle (red) approximated by circles of radius r_{disc} and other vehicle (blue) with shape- and uncertainty-ellipse corresponding to minimum occupancy probability p_e . Road boundaries $b_l(\theta)$ on the left and right $b_r(\theta)$ of reference path $(x^P(\theta), y^P(\theta))$.

E. Representation of Other Traffic Participants

For brevity we will refer to all traffic participants, such as vehicles, pedestrians, bicyclists, as vehicles. The shapes of other vehicles are approximated by a footprint encompassing ellipse of orientation ϕ with semi-major axes a_{shape} and b_{shape} in longitudinal and lateral direction of the obstacle respectively, cf. Fig. 3. For brevity, index i is omitted in this section. Consequently an analytical description of their occupied area is available, that will prove useful for describing collision states in closed analytic form. The evolution of their future trajectories are assumed to be known up to some uncertainty in the form of a posterior distribution and are parametrized by a mean trajectory $\mathbf{z}_{0:m-1}^i$ and uncertainty $\sigma_{0:m-1}$. In the more general case these should be supplied by an external inference framework. For our instantiation we supply a model of the growth of uncertainty

$$\sigma_{k+1} = \sigma_k + \sigma \Delta t_k, \quad (18)$$

of the vehicles position with uncertainty $\sigma_k = [\sigma_k^a, \sigma_k^b]^T$ at time k , and $\sigma = [\sigma^a, \sigma^b]^T$ the growth of uncertainty. The variances are approximated to be aligned with the vehicle's heading and thus the principle axis of the encompassing ellipse (cf. Fig. 3). The uncertainty growth in the lateral direction is bounded to a maximum value to take the high likelihood of vehicles staying in their current lanes into account.

The level-set lines of the Gaussian $\mathcal{N}(0, \text{diag}(\sigma_k))$ describing the position uncertainty of the other traffic participants at the level of p_ϵ form ellipses with coefficients

$$\begin{bmatrix} a_{\sigma_k} \\ b_{\sigma_k} \end{bmatrix} = \begin{bmatrix} \sigma_k^a \\ \sigma_k^b \end{bmatrix} \left(-2 \log(p_\epsilon 2\pi \sigma_k^a \sigma_k^b) \right)^{1/2}. \quad (19)$$

We can now use the axis alignment to the vehicle and directly add the coefficients to the semi-major axes to find the obstacle's ellipse with occupancy probability above the p_ϵ threshold.

The rectangular shape of the ego car is approximated by a set of discs of radius r_{ego} , cf. Fig. 3. It is necessary to employ discs instead of ellipses for the ego vehicle, since the ego vehicle and the other vehicles are not necessary axis aligned and the Minkowski sum can not be easily derived for non-axis aligned ellipses in closed form. The Minkowski sum of the ego car's discs and the previously derived occupancy-ellipse form analytic collision constraints

$$c_k^{\text{obst},i}(\mathbf{z}_k) = \left\| \begin{bmatrix} \Delta x_j \\ \Delta y_j \end{bmatrix}^T R(\phi)^T \begin{bmatrix} \frac{1}{a^2} & 0 \\ 0 & \frac{1}{b^2} \end{bmatrix} R(\phi) \begin{bmatrix} \Delta x_j \\ \Delta y_j \end{bmatrix} \right\|_{k,i} > 1, \quad \forall j \in \{1, \dots, 4\} \quad (20)$$

where $\Delta x, \Delta y$ are the distance of the ego vehicle's discs to the center of the obstacle i at time k . $R(\phi)$ is the rotation matrix corresponding to the obstacles heading, and

$$\begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a_{\text{shape}} + a_{\sigma_k} + r_{\text{disc}} \\ b_{\text{shape}} + b_{\sigma_k} + r_{\text{disc}} \end{bmatrix}. \quad (21)$$

the semi-major axes of the resulting constraint-ellipse. We now have an analytic constraint prohibiting collisions of probability higher than p_ϵ with other vehicles.

F. Minimal Intervention

It is our goal to follow the human input very closely and intervene only when deemed necessary. The minimal intervention cost term

$$J_h(\mathbf{z}_k, \mathbf{u}_k, \mathbf{u}_0^h) = \begin{bmatrix} u_k^a - a_0^h \\ \delta - \delta_0^h \end{bmatrix}^T K \begin{bmatrix} u_k^a - a_0^h \\ \delta - \delta_0^h \end{bmatrix} \quad (22)$$

penalizes the deviation of the system's state from the human driver commanded current controls $\mathbf{u}_0^h = [\delta_0^h, a_0^h]^T$, the steering angle δ^h and acceleration a^h at time t_k . In our setup we can only observe the driver steering angle δ^h and acceleration a^h , but not the steering speed $\dot{\delta}^h$. Nonetheless, the framework is general enough to also take the steering velocity as human input into account, if observable. The vehicle controls $\mathbf{u}_k$ remain steering *velocity* and acceleration.

G. Trajectory Cost

The trajectory cost contains the MPCC cost, Eq. (15), and additionally penalizes control inputs and yaw rate to create a smooth driving behavior and increase comfort. Weights R and A allow for different prioritization.

$$J_t(\mathbf{z}_k, \mathbf{u}_k, \theta_k) = J_{\text{MPCC}}(\mathbf{z}_k, \theta_k) + \mathbf{u}_k^T R \mathbf{u}_k + \dot{\phi}_k A \dot{\phi}_k \quad (23)$$

J_{MPCC} already encodes the penalization of the deviation from the reference path which results in a slight nudging behavior into a beneficial direction. It also takes the driver's goal of making progress along the road into account.

H. Optimization

We minimize the linear combination of the cost of intervention J_h (22) and trajectory cost J_t (23)

$$J(\mathbf{z}_k, \mathbf{u}_k, \theta_k, \mathbf{u}_0^h) = \beta \omega(t_k) J_h(\mathbf{z}_k, \mathbf{u}_k, \mathbf{u}_0^h) + (1 - \omega(t_k)) J_t(\mathbf{z}_k, \mathbf{u}_k, \theta_k) \quad (24)$$

weighted by β and an exponential decay function $w(t_k) = \exp(-\alpha t_k)$ to increase the impact of the human input in the short-term. We used a sharp drop-off, such that $w(0.5s) = 0.1$, and high values of β to make the system very responsive to human inputs but rely on J_t for steps further into the future. This strategy enables us to plan sufficiently well, without a prediction of driver intent or planned trajectory, since the planner's trajectory will *snap* into place shortly before the boundaries of constraints are met and is perceived as inactive to the human driver otherwise. We formulate the optimization problem with the aforementioned state-, dynamics-, path- and obstacle constraints and form the following constraint non-linear optimization problem:

$$\mathbf{u}_{0:m-1}^* = \arg \min_{\mathbf{u}_{0:m-1}} \sum_{k=0}^m J(\mathbf{z}_k, \mathbf{u}_k, \theta_k, \mathbf{u}_0^h) \Delta t_k \quad (25)$$

$$\text{s.t. } \mathbf{z}_{k+1} = f(\mathbf{z}_k, \mathbf{u}_k) \quad (26)$$

$$\theta_{k+1} = \theta_k + v_k \Delta t_k \quad (7)$$

$$\mathbf{z}_{\min} < \mathbf{z}_k < \mathbf{z}_{\max} \quad (27)$$

$$\mathbf{u}_{\min} < \mathbf{u}_k < \mathbf{u}_{\max} \quad (28)$$

$$\|\dot{\phi}_k\| < \dot{\phi}_{\max} \quad (29)$$

$$\|\phi_k - \phi^P(\theta_k)\| < \Delta \phi_{\max} \quad (17)$$

$$b_l(\theta_k) + w_{\max} \leq d(\mathbf{z}_k, \theta_k) \leq b_r(\theta_k) - w_{\max} \quad (16)$$

$$c_k^{\text{obstacle},i}(\mathbf{z}_k) > 1, \quad i = \{1, \dots, n\} \quad (21)$$

$$\forall k \in \{0, \dots, m\}$$

At initialization the path $(x^P(\theta), y^P(\theta))$ and boundaries $b_l(\theta)$ and $b_r(\theta)$ are given by the road and static obstacles, cf. Fig. 4. At the beginning of each control loop the initial states $\mathbf{z}_0, \theta_0$, human control input $\mathbf{u}_0^h$, and predictions of other traffic participants $\mathbf{z}_{0:m}^i, \sigma_{0:m}^i$ are provided to the NMPC. After solving Eq. (25) the optimal control $\mathbf{u}_0^*$ is executed by the system. The optimization problem is solved by a Primal-Dual Interior Point solver generated by FORCES Pro [20].

Fig. 4. Control scheme of the NMPC

V. RESULTS

We evaluate the capabilities of our approach in a variety of simulated scenarios. The human driver controls a physical steering wheel and pedals which generate the desired steering angle δ_0^h and acceleration a_0^h . The inputs are then processed in the MPC formulation to guarantee safe motion. The reference path and the road boundaries b_l and b_r are designed to fit the road network.

A. Sharp Turn

In this scenario, cf. Fig. 5(a), the vehicle enters a sharp left turn on a race track. The current human inputs would cause the vehicle to quickly leave the road at high speed, as shown by the red line. The controller brakes the vehicle to a safe speed complying with the yaw-rate constraint, then accelerates at the exit of the turn to maximize progress, while always respecting the roads limits.

The planned trajectory shows similarities to a racing line during high-speed cornering. This behavior shows the advantage of longitudinal and lateral control; without deceleration the vehicle would not have been able to complete the turn shown by the red line in Fig. 5(a). The plan maintains a smooth acceleration profile during the turn, cf. Fig. 5(b).

B. Left Turn Intersection and Merging Into Traffic

In this challenging left turn cross traffic scenario the initial position and velocity profiles of all other traffic participants are randomly generated. To increase the time horizon of the planner without sacrificing computation, we adopt a variable step size approach. For sufficient temporal resolution in the short term the first 10 steps are spaced by $\Delta t_k = 0.1s$ and $\Delta t_k = 0.2s$ for the remaining 40 steps, resulting in a planning horizon of nearly 9s with all computations done in real-time, cf. Sec.V

We have evaluated the method in a set of 50 randomly generated scenarios without any resulting collisions despite deliberately giving the system unsafe inputs which would have resulted in crashes without the proposed Parallel Autonomy system. Furthermore, the system did not cause any collisions which would not have happened without it.

We present two representative examples, for two different human driving styles. In the first case, cf. Fig. 6, an aggressive driver nearly collides with the right road boundary even before entering the intersection ①. Then, the driver tries to accelerate into the intersection ②, although other vehicles are just passing. Our

(a) MPC plan of vehicle, where vehicle poses are 0.1s apart for a horizon of 5s. Red line shows the constant steering angle and acceleration propagation of the human input for a horizon of 2s.

(b) The limit of yaw rate (1rad/s) (left) causes the vehicle to decelerate (right) and accelerate again at the end of the turn.

Fig. 5. Sharp Turn: Output of MPCC plans to decelerate into the sharp corner to comply with a yaw-rate constraint of 1rad/s.

system brings the ego vehicle to a full stop, lets the other vehicles pass and then proceeds by letting the driver merge into the traffic when a large enough gap appears. At ③ the driver approaches a preceding vehicle with high relative speed and tries to collide by accelerating even further. Our system brakes the ego vehicle and allows an overtaking maneuver once the oncoming traffic has passed. At ④ the driver erratically tries to break through the right road boundary, which is prohibited by our system. In all these cases the system can guard the human driver from actually causing any harm to himself and others.

The opposite spectrum of how our method reacts is shown in Fig. 7: A fairly good and calm driver experiences the same previous scenario. We observe that if the inputs from the human driver are deemed safe, barely any difference between human and system inputs occurs. The system thus minimizes intervention if no critical situations occur. Since steering the vehicle with steering wheel and pedals in simulation is not an easy task, due to the lack of feedback, the human driver did not break sufficiently at ② and misses to counter steer during a lane change maneuver ④. Notice how the system allows the driver to stay stationary at the intersection longer than necessary for safety. We can see that the trajectory cost J_t , which includes a path-progress term, does not cause the vehicle to start driving if the human driver does not intend to. The short term impact of the minimal intervention cost J_h always dominates the trajectory optimization if safety can be assured.

C. Impact of Uncertainty

Taking the uncertainty in the prediction of other vehicles into account is important, since future states can deviate substantially from the expectation. In the case of neglecting uncertainties the planned behavior can be more aggressive and is given larger leeway in the constraints. See Fig. 8-bottom, where the vehicle is

Fig. 6. *Aggressive left turn with traffic*: The system's steering angle and acceleration are displayed in blue, the human input in red. Snapshots of the current scenes at specific time-stamps are displayed above the acceleration and steering plots: The ego vehicle in red, the MPC planned path in blue. All other vehicles in black. An aggressive driver causes multiple critical situations where the system is forced to intervene to large amounts to keep the vehicle in a safe state. Large deviation from the driver's desired acceleration and steering wheel angle to the actual system output are observable. E.g. collision at time (2) is prohibited by strong braking.

Fig. 7. *Normal left turn with traffic*: System output stays close to the desired human acceleration and steering wheel angle. An exception appears at (4) where the driver is not counter steering enough to prohibit a predicted collision with the left road boundary.

allowed to merge into the lane in front of a second vehicle. Taking future obstacles' uncertainty growth into account, cf. Fig. 8-top, results in more conservative behavior and the ego vehicle is prohibited from merging.

D. Computation Time

The NMPC solve-times collected during several runs are displayed in Fig. 9. Results were computed on an off-the-shelf Intel Core i5-4200U mobile CPU @ 1.6 GHz, 2.6 GHz Turbo Boost, and 6 GB RAM. We observed a strong influence of the complexity of the scenario on the computation time. In the case of no dynamic obstacles we saw solve-times of less than 30ms

Fig. 8. Comparison of NMPC plans with uncertainty estimate (top) and without (bottom) shown by the ellipses representing their occupancy probability threshold. Predicted future states are shown in fading colors 0.4s apart over a horizon of 9s.

Fig. 9. Computation times for the different human in the loop scenarios and varying number of other traffic participants, and high uncertainty. Finally the case of AI only, without the human in the loop.

even for a challenging race track with many tight turns, forcing the MPC to intervene and decelerate due to turn-rate constraints. In cases where the system needs to nudge into tight gaps while simultaneously deciding whether a subsequent overtaking maneuver is feasible, computation times can reach up to 65ms in exceptional cases. Our system was able to reach the goal replanning frequency of 10Hz at all times.

VI. CONCLUSION

In this work we presented a receding horizon planner that minimizes deviation from the human input while ensuring safety according to our proposed general Parallel Autonomy control framework. We have shown the increased functionality compared to other approaches in complex and more realistic driving scenarios.

Future work will include evaluation on a real vehicle platform as well as a more involved system model derived via an identification step including combined slip and load-transfers. Further tests will also enclose an inference framework to gain more elaborate predictions of other traffic participants. The presented framework may be applied to more general scenarios including a larger variety of dynamic obstacles such as pedestrians, bicycles, trucks, as well as a larger variety of environments including stop-signs and traffic lights by small adaption of our set of constraints.

Furthermore, the proposed receding horizon planner also applies to fully autonomous vehicles if the minimal intervention cost is excluded and future experiments will show the functionality.

The widespread deployment of our method in vehicle systems would help to reduce the massive number of vehicular injuries and fatalities, as well as provide a safe pathway towards the development of fully autonomous vehicles.

REFERENCES

- [1] ASIRT, "Association for Safe International Road Travel 2016."
- [2] National Safety Council (NSC), "Motor Vehicle Fatality Estimates 2012-2015."
- [3] NHTSA, "Traffic Safety Facts 2015."
- [4] M. Nohmi, N. Fujikura, C. Ueda, and E. Toyota, "Automatic braking or acceleration control system for a vehicle," Jan. 3 1978, US Patent 4,066,230.
- [5] E. Dagan, O. Mano, G. P. Stein, and A. Shashua, "Forward collision warning with a single camera," in *IEEE Intelligent Vehicles Symposium, 2004*, June 2004, pp. 37–42.
- [6] M. Forghani, J. M. McNew, D. Hoehener, and D. D. Vecchio, "Design of driver-assist systems under probabilistic safety specifications near stop signs," *IEEE Transactions on Automation Science and Engineering*, vol. 13, no. 1, pp. 43–53, Jan 2016.
- [7] D. Hoehener, G. Huang, and D. D. Vecchio, "Design of a lane departure driver-assist system under safety specifications," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, Dec 2016, pp. 2468–2474.
- [8] T. Fraichard and H. Asama, "Inevitable collision states. A step towards safer robots?" *Advanced Robotics*, vol. 18, no. 10, pp. 1001–1024, 2004.
- [9] A. Bautin, L. Martinez-Gomez, and T. Fraichard, "Inevitable Collision States: A probabilistic perspective," in *2010 IEEE International Conference on Robotics and Automation*, May 2010, pp. 4022–4027.
- [10] D. Althoff, M. Althoff, D. Wollherr, and M. Buss, "Probabilistic collision state checker for crowded environments," in *2010 IEEE International Conference on Robotics and Automation*, May 2010, pp. 1492–1498.
- [11] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, "A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games," *IEEE Transactions on Automatic Control*, vol. 50, no. 7, pp. 947–957, July 2005.
- [12] V. A. Shia, Y. Gao, R. Vasudevan, K. D. Campbell, T. Lin, F. Borrelli, and R. Bajcsy, "Semiautonomous vehicular control using driver modeling," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 6, pp. 2696–2709, Dec 2014.
- [13] Y. Gao, A. Gray, A. Carvalho, H. E. Tseng, and F. Borrelli, "Robust nonlinear predictive control for semiautonomous ground vehicles," in *2014 American Control Conference*, June 2014, pp. 4913–4918.
- [14] S. M. Erlien, S. Fujita, and J. C. Gerdes, "Shared steering control using safe envelopes for obstacle avoidance and vehicle stability," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 2, pp. 441–451, Feb 2016.
- [15] J. Alonso-Mora, P. Gohl, S. Watson, R. Siegwart, and P. Beardsley, "Shared control of autonomous vehicles based on velocity space optimization," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 1639–1645.
- [16] A. Gray, Y. Gao, T. Lin, J. K. Hedrick, H. E. Tseng, and F. Borrelli, "Predictive control for agile semi-autonomous ground vehicles using motion primitives," in *2012 American Control Conference (ACC)*, June 2012, pp. 4239–4244.
- [17] S. J. Anderson, S. B. Karumanchi, and K. Iagnemma, "Constraint-based planning and control for safe, semi-autonomous operation of vehicles," in *2012 IEEE Intelligent Vehicles Symposium*, June 2012, pp. 383–388.
- [18] B. Paden, M. p, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, March 2016.
- [19] S. J. Anderson, S. C. Peters, T. E. Pilutti, and K. Iagnemma, "An optimal-control-based framework for trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles in hazard avoidance scenarios," *International Journal of Vehicle Autonomous Systems*, vol. 8, no. 2-4, pp. 190–216, 2010.
- [20] A. Domahidi and J. Jerez, "FORCES Pro," embotech GmbH (<http://embotech.com/FORCES-Pro>), July 2014.
- [21] D. Lam, C. Manzie, and M. C. Good, "Model predictive contouring control for biaxial systems," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 2, pp. 552–559, March 2013.
- [22] T. Faulwasser, B. Kern, and R. Findeisen, "Model predictive path-following for constrained nonlinear systems," in *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, Dec 2009, pp. 8642–8647.
- [23] 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.