The International Journal of
Robotics Research
31(3) 346–359

DOI: 10.1177/0278364911423042
ijr.sagepub.com

Optimal trajectories for time-critical street scenarios using discretized terminal manifolds

Moritz Werling[1,2], Sören Kammel[3], Julius Ziegler[4] and Lutz Gröll[5]

Abstract

This paper deals with the trajectory generation problem faced by an autonomous vehicle in moving traffic. Being given the predicted motion of the traffic flow, the proposed semi-reactive planning strategy realizes all required long-term maneuver tasks (lane-changing, merging, distance-keeping, velocity-keeping, precise stopping, etc.) while providing short-term collision avoidance. The key to comfortable, human-like as well as physically feasible trajectories is the combined optimization of the lateral and longitudinal movements in street-relative coordinates with carefully chosen cost functionals and terminal state sets (manifolds). The performance of the approach is demonstrated in simulated traffic scenarios.

Keywords

Autonomous vehicle, motion planning, moving obstacles, optimal control

1. Introduction

The past three decades have witnessed an ambitious research effort in the area of automated driving. This has led to a remarkable enhancement in terms of handling complex situations robustly. The further autonomous vehicles advance towards realistic road traffic the more often they face dynamic, time-critical street scenarios, such as merging into fast traffic flow, to pass opposing traffic, or to avoid other moving vehicles. Under simplified conditions, such as during the 2007 *DARPA Urban Challenge*,[1] these dynamic scenarios can be tackled with path-based planning strategies in combination with fairly simple heuristics and conservative estimates (Werling et al. 2008). However, in rushing nose-to-tail traffic, these approaches quickly reach their limits resulting in poor performance or even accidents (Fletcher et al. 2008). This is where trajectory-based planning concepts come into play, which explicitly account for the time t.

1.1. Related work

The task of creating time parameterized trajectories, and accounting for kinematic constraints at the same time is referred to as *kinodynamic planning*. Several methods have been proposed, which find a *global* trajectory connecting the initial state with an *exactly defined* goal state (see LaValle and Kuffner (2001) for a survey). They often rely on discrete geometric structures, including rapidly exploring random trees (RRT) (LaValle 1998) and state lattices (Pivtoraiko and Kelly 2005). Both have been applied successfully in full-sized autonomous vehicles (Kuwata et al. 2008; Likhachev and Ferguson 2009). While being most suitable for combinatorially difficult problems encountered in unstructured environments like parking lots, these methods typically cannot quickly consider *alternative* goal states. This, however, may be required for an evasive maneuver. Also, these combinatorial methods are typically computationally expensive, which hampers short planning cycles. Both the capability to consider a set of alternative goal states and high replanning frequencies are important preconditions to adopt a motion planning method for high-speed obstacle avoidance.

We therefore embark on the strategy proposed in Howard and Kelly (2007) and Montemerlo et al. (2008). It takes advantage of the structure in the road environment by deliberating *multiple* final states (cf. "ego-graph" (Lacaze et al.

[1]Department of Applied Computer Science and Automation (AIA), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
[2]BMW Group Research and Technology, Munich, Germany
[3]Robert Bosch LLC Research and Technology Center, Palo Alto, California, USA
[4]Department of Measurement and Control (MRT), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
[5]Department of Applied Computer Science (IAI), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany

Corresponding author:
Moritz Werling, BMW Group Research and Technology, Hanauerstr. 46, 80992 Munich, Germany
Email: moritz.werling@bmw.de

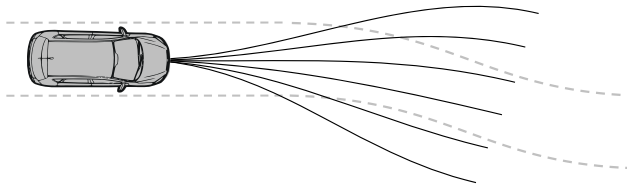

Fig. 1. Fan-shaped trajectory alternatives offsetting the vehicle from the street center for obstacle avoidance.

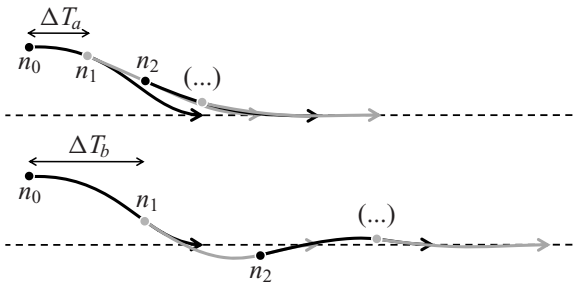

Fig. 2. Varying transient behavior caused by different cycle times ΔT_a and ΔT_b: (top) short cycle time with tolerable transient; (bottom) long cycle time causing overshoots, the n_i denote the starting points of subsequent planning steps.

1998), see Figure 1). In combination with short replanning cycles (e. g. 100 ms) this leads to a reactive layer, which is highly responsive to traffic changes.

With the permanent replanning, however, a critical effect comes along, if we incorporate heuristics or suboptimality. More precisely, the discrepancy between consecutive plans can easily result in overshoots, oscillations, or even instabilities of the vehicle movement, which is especially critical at high speeds. Furthermore, the emergent trajectory is significantly affected by the planning cycle time, as outlined in Figure 2.

However, a *temporal consistency* (TC), as we refer to what follows, cannot always be accomplished in practice. On the one hand, the previous assumptions about the future obstacle trajectories might prove wrong, so that course correction according to the new sensor data is necessary. On the other, computing time, and therefore the optimization horizon, is finite, so that in each step the new information will be incorporated.[2] As for the proposed strategy in this paper, TC is only assured whenever it is not practically impeded for the above reasons.

According to *Bellman's Principle of Optimality*, an optimal policy implies TC (Bellman 1954). This fact, rather than a specific objective function, is the reason for the upcoming optimal control approach.

Based on experience, dense highway traffic also necessitates the *combined* optimization of the lateral (steering) and longitudinal movement (gas/brake), as the vehicular movement is in equal measures influenced by the pedal and steering usage. That is, full braking can often be avoided by well-dosed steering action as well as large steering amplitudes by a cautious deceleration. This combined optimization also sets us apart from the existing reactive solutions such as Kelly and Nagy (2003), which only utilize the lateral movement for obstacle avoidance.

On top of the aforementioned features, our method (a first draft has been presented in Werling et al. (2010b)) is characterized by unobtrusive, comfortable maneuvers at a wide speed range. It also provides for a *unified* treatment of all required operation modes (velocity-keeping, distance-keeping, etc.) as well as a simple and tight integration into the behavioral layer with intuitive, orthogonal parameter tuning, which we found to be extremely useful in practice.

1.2. Paper outline

Examining the different aspects of the proposed algorithm, we organize the remainder of the paper as follows. After explaining the rationale behind the street-relative coordinate choice in Section 2, all optimization restrictions (vehicular physics and obstacles) are initially disregarded. This allows in the first part of Section 3 for a generic, unified treatment of the lateral and longitudinal movement by optimal control theory leading to fast computable closed-form solutions. In the second part, the derived formulas are applied to the original, restricted optimization problem. Section 4 then describes the final application. Moreover, the brief description of the prototypical setup of a full-size road vehicle in Section 5 is followed by the illustration of the algorithm's performance in various simulated traffic scenarios. The results as well as additional aspects of the proposed method are then discussed in Section 6. Finally, a short summary and an outlook are given in Section 7.

2. Problem formulation in the Frenet coordinates

To achieve a human-like driving experience we need to understand the coordinate system humans plan in when they drive. Consider the lane change situation shown in Figure 3. Optimizing for comfort, time and energy the lane change (a) is superior to (b). Yet humans will often maneuver as shown in (b) which executes strategy (c). The human is planning the vehicle's lateral movement relative to the lanes rather than to the absolute ground.

Imitating this approach, the trajectory generation problem is formulated in the so-called Frenet frame $[\boldsymbol{n}_c, \boldsymbol{t}_c]$ of the street as shown in Figure 4. The timed offset to the lane center is denoted by $d(t)$ and the lane is represented by its center line (solid curve).

In addition to that it would now stand to reason to use the movement along $\boldsymbol{t}_x$ such as the covered arc length of the vehicle as the remaining second degree of freedom. Instead, we choose the covered arc length $s(t)$ of the frame's root

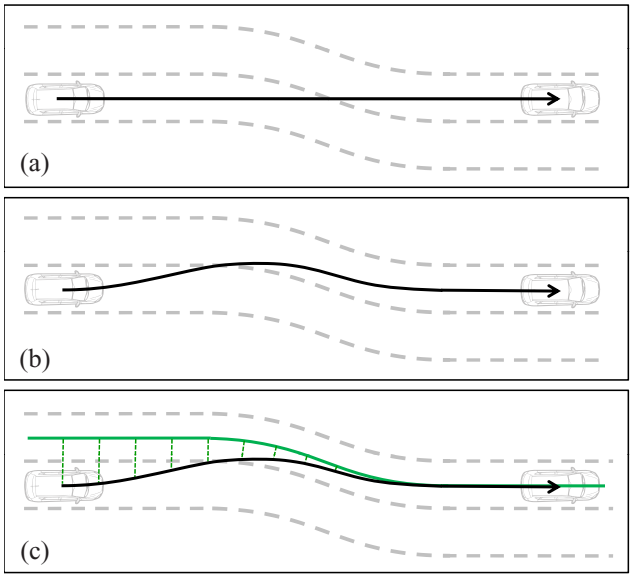

Fig. 3. Possible lane change maneuvers (a) and (b) as well as the visualization (c) of the strategy behind (b).

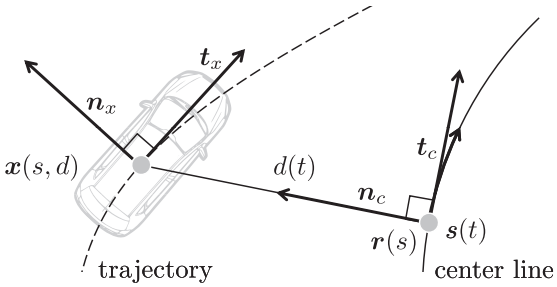

Fig. 4. Trajectory representation in the Frenet coordinates of the street.

point $\boldsymbol{r}(s)$ along the center line allowing the fast computable closed-form reparameterization

$$\boldsymbol{x}(s(t),d(t)) = \boldsymbol{r}(s(t)) + d(t)\,\boldsymbol{n}_c(s(t)) \tag{1}$$

of the planned trajectory $\boldsymbol{x}(t)$.

Aiming at an optimal policy, we next define a suitable optimality criterion. In order to allow for modularity among the different operation modes, such as velocity and distance-keeping, we assume that the trajectory costs J to minimize may be separated into a lateral and a longitudinal component, J_d and J_s, according to the weighted sum

$$J[d,s] = J_d[d] + k_s J_s[s], \quad k_s > 0. \tag{2}$$

This also allows for the unified treatment of the trajectory's lateral and longitudinal movement in the next section.

3. Optimal control formulation

In order to unify the treatment of the lateral and longitudinal movement $d(t)$ and $s(t)$, we consider them as the outputs $d(t) = \xi_1(t)$ and $s(t) = \xi_1(t)$ *respectively* of *separate* integrator systems of the form

$$\dot{\boldsymbol{\xi}} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \boldsymbol{\xi} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} u =: \boldsymbol{f}(\boldsymbol{\xi}, u) \tag{3}$$

with $\boldsymbol{\xi}^{\mathrm{T}} = [\xi_1, \xi_2, \xi_3]$ and subject to no restrictions. With the input $u(t) = \dddot{\xi}_1(t)$ being the lateral/longitudinal so-called jerk $\dddot{d}(t)$ / $\dddot{s}(t)$, we can define the cost functional

$$J_\xi := \int_0^\tau f_0(u)\,\mathrm{d}t + (h(\boldsymbol{\xi}(t),t))_\tau, \quad f_0(u) := \frac{1}{2}u^2(t) \tag{4}$$

with yet both unspecified terminal costs $(h(\boldsymbol{\xi}(t),t))_\tau$ and terminal time τ. We propose the following theorem.

Theorem 1. *The unconstrained movement of $\xi_1(t)$ that transfers the system* (3) *from the initial state $\boldsymbol{\xi}(0) = \boldsymbol{\xi}_0$ to a given final state $\boldsymbol{\xi}(\tau) = \boldsymbol{\xi}_\tau$ minimizing* (4) *is in the set of quintic polynomials.*

Proof. Differentiating the Lagrangian

$$\begin{aligned}\mathcal{L} :=& f_0 - \boldsymbol{\psi}^{\mathrm{T}}[\boldsymbol{f} - \dot{\boldsymbol{\xi}}] \\ =& \frac{1}{2}u^2 + \psi_1[\dot{\xi}_1 - \xi_2] + \psi_2[\dot{\xi}_2 - \xi_3] + \psi_3[\dot{\xi}_3 - u]\end{aligned}$$

with respect to the input, $\boldsymbol{\psi}$ being the Lagrange multipliers, reduces the control equation $\frac{\partial \mathcal{L}}{\partial u} = 0$ to

$$u = \psi_3. \tag{5}$$

With this, the Euler–Lagrange equation $\frac{\partial \mathcal{L}}{\partial \boldsymbol{\xi}} - \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\partial \mathcal{L}}{\partial \dot{\boldsymbol{\xi}}}\right) = 0$ yields

$$0 - \dot{\psi}_1 = 0, \tag{6a}$$
$$-\psi_1 - \dot{\psi}_2 = 0, \tag{6b}$$
$$-\psi_2 - \dot{\psi}_3 = 0. \tag{6c}$$

From (6a) it follows that $\psi_1 = \text{const.}$ and with (6b), (6c) and (3) we get

$$\begin{bmatrix} \xi_1 \\ \dot{\xi}_2 \\ \ddot{\xi}_3 \\ \dddot{\psi}_3 \\ \ddot{\psi}_2 \\ \dot{\psi}_1 \end{bmatrix} = \begin{bmatrix} 1 & t & t^2 & t^3 & t^4 & t^5 \\ 0 & 1 & 2t & 3t^2 & 4t^3 & 5t^4 \\ 0 & 0 & 2 & 6t & 12t^2 & 20t^3 \\ 0 & 0 & 0 & 6 & 24t & 60t^2 \\ 0 & 0 & 0 & 0 & -24 & -120t \\ 0 & 0 & 0 & 0 & 0 & 120 \end{bmatrix} \begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ c_3 \\ c_4 \\ c_5 \end{bmatrix} \tag{7}$$

with $c_0, c_1, \ldots, c_5 \in \mathbb{R}$. □

Similarly, we propose the following theorem.

Theorem 2. *The unconstrained movement of* $\xi_1(t)$ *minimizing* (4) *that transfers the system* (3) *from the initial state* $\boldsymbol{\xi}(0)=\boldsymbol{\xi}_0$ *to a given final state* $\boldsymbol{\xi}(\tau)=\boldsymbol{\xi}_\tau$*, but with an arbitrary end position* $\xi_1(\tau)$*, is in the set of quartic polynomials if* $h(\boldsymbol{\xi}(t),t)$ *is independent of* ξ_1.

Proof. The proof is identical to the previous one, except for the additional transversality equation $\psi_1=-\frac{\partial h}{\partial \xi_1}=0$ leading to $c_5=0$. □

Due to the three initial and final states, the coefficients $[\boldsymbol{c}_{012}^{\mathrm{T}},\boldsymbol{c}_{345}^{\mathrm{T}}] := [c_0,c_1,c_2,c_3,c_4,c_5]$ of the quintic state trajectory

$$\boldsymbol{\xi}(t)=\mathbf{M}_1(t)\,\boldsymbol{c}_{012}+\mathbf{M}_2(t)\,\boldsymbol{c}_{345} \tag{8}$$

with

$$\mathbf{M}_1(t):=\begin{bmatrix}1 & t & t^2\\ 0 & 1 & 2t\\ 0 & 0 & 2\end{bmatrix},$$
$$\mathbf{M}_2(t):=\begin{bmatrix}t^3 & t^4 & t^5\\ 3t^2 & 4t^3 & 5t^4\\ 6t & 12t^2 & 20t^3\end{bmatrix}$$

are for $\tau>0$ exactly defined by

$$\boldsymbol{c}_{012}=\mathbf{M}_1(0)^{-1}\,\boldsymbol{\xi}_0=:\boldsymbol{q}_{012}(\xi_1(0),\xi_2(0),\xi_3(0)), \tag{9}$$
$$\begin{aligned}\boldsymbol{c}_{345}&=\mathbf{M}_2(\tau)^{-1}\,[\boldsymbol{\xi}_\tau-\mathbf{M}_1(\tau)\,\boldsymbol{c}_{012}]\\ &=:\boldsymbol{q}_{345}(\xi_1(\tau),\xi_2(\tau),\xi_3(\tau)).\end{aligned} \tag{10}$$

The coefficients of the quartic polynomial can be analogously determined by (9) and

$$\boldsymbol{c}_{34}=\boldsymbol{q}_{34}(\xi_2(\tau),\xi_3(\tau)).$$

With this we only reduced the variational calculus problem to an optimization problem since the terminal time τ as well as the final states $[\xi_1(\tau),\xi_2(\tau),\xi_3(\tau)]$ and $[\xi_2(\tau),\xi_3(\tau)]$ still have to be suitably selected. This will be done in the next section.

3.1. Introducing a continuous terminal manifold

When optimizing the terminal state $\boldsymbol{\xi}_\tau$ of the lateral and longitudinal movement, the target application narrows a priori the set of reasonable solutions. For one, also on a partially blocked road, the vehicle should generally progress along it and not crosswise. With this in mind, we define, first for the quintic polynomials, a terminal manifold given by

$$(z(\boldsymbol{\xi}(t),t))_\tau:=\left(\begin{bmatrix}\xi_2(t)-\dot{\xi}_{\mathrm{ref}}(t)\\ \xi_3(t)-\ddot{\xi}_{\mathrm{ref}}(t)\end{bmatrix}\right)_\tau=\mathbf{0}, \tag{11}$$

constraining ξ_1's first and second derivative $\xi_2(t)$ and $\xi_3(t)$ at time τ to be identical to those of a reference trajectory $\xi_{\mathrm{ref}}(t)$ (cf. the dashed lines in Figure 5).

As for the choice of the terminal time, reaching the final state early might lead to uncomfortable, energetically wasteful actions, whereas a late arrival implies lagging movements. Since these issues are also strongly coupled with the final state on the terminal manifold we seek to find the best trade-off by defining the terminal costs in (4) to be

$$(h(\boldsymbol{\xi}(t),t))_\tau:=\left(k_\tau t+\frac{1}{2}k_{\xi_1}[\xi_1(t)-\xi_{\mathrm{ref}}(t)]^2\right)_\tau \tag{12}$$

with $k_\tau,k_{\xi_1}>0$, which penalize both slow convergence and final deviations from the reference trajectory. As Theorem 1 is independent of the terminal costs, it suffices to find the quintic polynomial's pair $[\xi_1(\tau),\tau]$ on the terminal manifold, that minimizes the total costs (4) with (12).

Remark 1. *The following steps, applying to quintic polynomials only, are not required for implementation but provide both an insight into the proposed algorithm and useful estimates for the choice of parameters. Analog calculations can be carried out for quartic polynomials.*

Following optimal control theory, with

$$\left(\frac{\partial z}{\partial \boldsymbol{\xi}}\right)_\tau=\begin{bmatrix}0 & 1 & 0\\ 0 & 0 & 1\end{bmatrix}$$

and (12) the transversality equation reads

$$\left(\frac{\partial h}{\partial \boldsymbol{\xi}}\right)_\tau+\boldsymbol{\psi}(\tau)-\left(\frac{\partial z}{\partial \boldsymbol{\xi}}\right)_\tau^{\mathrm{T}}\boldsymbol{\mu}=$$
$$\begin{bmatrix}k_{\xi_1}[\xi_1(\tau)-\xi_{\mathrm{ref}}(\tau)]\\ 0\\ 0\end{bmatrix}+\begin{bmatrix}\psi_1(\tau)\\ \psi_2(\tau)\\ \psi_3(\tau)\end{bmatrix}-\begin{bmatrix}0 & 0\\ 1 & 0\\ 0 & 1\end{bmatrix}\begin{bmatrix}\mu_1\\ \mu_2\end{bmatrix}=0.$$

The first row can be rewritten as

$$[\xi_1(\tau)-\xi_{\mathrm{ref}}(\tau)]=-\frac{1}{k_{\xi_1}}\psi_1=-\frac{120}{k_{\xi_1}}c_5. \tag{13}$$

This combined with (11), solving for

$$\boldsymbol{\xi}_{\mathrm{ref}}:=[\xi_{\mathrm{ref}},\dot{\xi}_{\mathrm{ref}},\ddot{\xi}_{\mathrm{ref}}]^{\mathrm{T}},$$

as well as using (8), gives

$$\boldsymbol{\xi}_{\mathrm{ref}}(\tau)=\mathbf{M}_1(\tau)\,\boldsymbol{c}_{012}+\mathbf{M}_3(\tau)\,\boldsymbol{c}_{345}$$
$$\text{with}\quad \mathbf{M}_3(\tau):=\mathbf{M}_2(\tau)+\begin{bmatrix}0 & 0 & \frac{120}{k_{\xi_1}}\\ 0 & 0 & 0\\ 0 & 0 & 0\end{bmatrix}.$$

As $\tau,k_{\xi,1}>0$, this can be solved for the optimal parameters $\boldsymbol{c}_{345}$ yielding

$$\boldsymbol{c}_{345}(\tau)=\mathbf{M}_3^{-1}(\tau)\,[\boldsymbol{\xi}_{\mathrm{ref}}(\tau)-\mathbf{M}_1(\tau)\,\boldsymbol{c}_{012}] \tag{14}$$

with $\boldsymbol{c}_{012}$ given by (9). Notice, that this is equivalent to determining the optimal final state $\boldsymbol{\xi}_\tau(\tau)$.

Optimizing also the time τ to the terminal manifold, the equation

$$f_0(u)-\boldsymbol{\psi}(\tau)^{\mathrm{T}}\boldsymbol{f}(\tau)+\left(\frac{\partial h}{\partial t}\right)_\tau=0$$

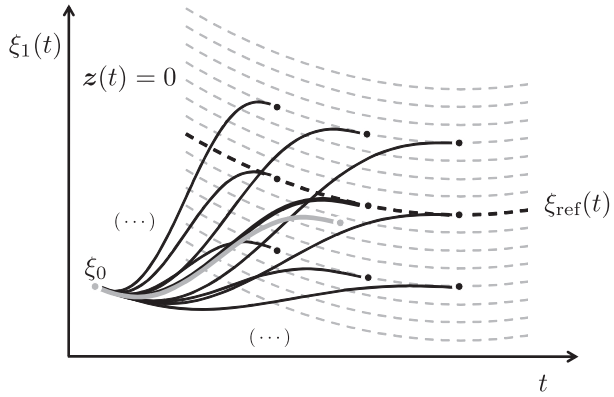

Fig. 5. Optimal trajectory (thick, gray) with end point on the continuous terminal manifold (dashed in gray) as well as the optimal trajectory (thick, black) and selected alternatives (thin, black) to the discretized manifold.

holds. With

$$\left(\frac{\partial h}{\partial t}\right)_\tau = (k_\tau + k_{\xi_1}[\xi_1(t)-\xi_{\mathrm{ref}}(t)]\underbrace{[\dot{\xi}_1(t)-\dot{\xi}_{\mathrm{ref}}(t)]}_{=z_1=0})_\tau = k_\tau,$$

it can be written with (5) as

$$\frac{1}{2}u^2(\tau)-\psi_1(\tau)\xi_2(\tau)-\psi_2(\tau)\xi_3(\tau)-\psi_3(\tau)u(\tau)+k_\tau = -\frac{1}{2}\psi_3^2(\tau)-\psi_1(\tau)\dot{\xi}_{\mathrm{ref}}(\tau)-\psi_2(\tau)\ddot{\xi}_{\mathrm{ref}}(\tau)+k_\tau = 0, \quad (15)$$

whereas $\psi_1(\tau)$, $\psi_2(\tau)$, and $\psi_3(\tau)$ are given by (7) and (14). For arbitrary functions $\xi_{\mathrm{ref}}(t)$ the optimal time τ (cf. Figure 5) can only be determined by solving (15) numerically. As (15) is only a necessary condition, the optimal trajectory might need to be singled out among multiple solutions by an explicit evaluation of the cost functional. This approach brings us to the actually implemented strategy in the next section.

3.2. Numerical optimization over a discretized terminal manifold

Unfortunately, the original optimization problem is much more complicated due to the so far unconsidered physical and collision-related constraints. The latter preclude the identification of an optimal function class[3] such as in Theorems 1 and 2. In order to avoid an expensive numerical optimization of the trajectory shape, we propose the following simple strategy.

Optimization heuristic. *Select the best solution in the function space of the optimal solution to the unconstrained problem that meets the constraints.*

As the constraints are predominantly not active in practice (this is the task of the behavior layer, cf. Section 6), the optimal solution is obtained – and therefore temporal consistency – in the majority of cases. Otherwise, still intuitive trajectories are generated as will be shown in Section 5.

Even with this simple strategy, there are yet infinitely many quintic/quartic polynomials leading to the terminal manifold, which have to be tested in each planning cycle for compliance with the constraints (cf. Section 4.4). The remedy for this problem is a conscious reduction of the permitted solutions by substituting the terminal manifold for a discrete set of its terminal states, similar to the generation of a lattice. We now only allow the trajectory to arrive at certain points in absolute[4] time as well as with certain discrete distances to the reference trajectory $\xi_{\mathrm{ref}}(t)$, referred to as *target points* below. Consequently, the admissible polynomials form an entire fan-shaped trajectory *set* evenly covering the maneuver space as shown in Figure 5 and we know the following from Bellman's Principle.

Corollary 3. *Selecting the best trajectory by an explicit evaluation of the cost functional provides with inactive constraints an optimal policy for the discretized terminal manifold – and therefore temporal consistence.*

It goes without saying that in practice only a finite number of trajectories to the terminal manifold can be inspected. This does not pose a problem as the chosen terminal costs (12) prefer trajectories reaching the terminal manifold in the near future so that only those need to be evaluated. For implementation, the required manifold horizon can be estimated by (15).

Remark 2. *Substituting the last component of* (14)*, namely*

$$c_5 = -[12[\xi_1(0)-\xi_{\mathrm{ref}}(\tau)]+6\tau[\xi_2(0)+\dot{\xi}_{\mathrm{ref}}(\tau)] + \tau^2[\xi_3(0)-\ddot{\xi}_{\mathrm{ref}}(\tau)]]\frac{k_{\xi_1}}{1440+2\tau^5 k_{\xi_1}},$$

in (13) *for the special case* $\xi_2(0)=\xi_3(0)=\dot{\xi}_{\mathrm{ref}}(\tau)=\ddot{\xi}_{\mathrm{ref}}(\tau)=0$ *yields the relation*

$$[\xi_1(\tau)-\xi_{\mathrm{ref}}(\tau)] = \frac{1}{1+\frac{1}{720}\tau^5 k_{\xi_1}}[\xi_1(0)+\xi_{\mathrm{ref}}(0)]. \quad (16)$$

Obviously, the optimal solution onto the continuous terminal manifold is not right on the reference trajectory at the end (cf. Figure 5) but offset proportionally to the initial distance. The discretized terminal manifold now inures to our benefit, as with an estimated upper bound on $|[\xi_1(0)+\xi_{\mathrm{ref}}(0)]|$ *(e. g. one lane width) parameter* k_{ξ_1} *can be chosen with* (15) *and* (16)*, so that the optimal trajectory 'latches' onto the reference (cf. Figures 5–9).*

Remark 3. *The proposed optimization can only provide the aspired temporal consistency up to reaching the terminal manifold, as the problem was only defined until then. Afterwards, however, as stated in the previous remark, with inactive constraints (that is no obstacles) the trajectory leads to the terminal state with* $h=0$ *and stays there consistently.*

This completes our calculations for the unconstrained integrator system required later in the actual implementation. For the original trajectory planning problem, we still need to come up with problem orientated choices for $\xi_{\text{ref}}(t)$ and combine the different lateral and longitudinal movements $d(t)$ and $s(t)$ before selecting the final trajectory $\boldsymbol{x}(t)$.

Remark 4. *The suboptimality induced by the discretization in the manifold is completely negligible in practice, as the minimization criterion is only an auxiliary construct. However, the target points on the manifold cannot be arbitrarily sparse in the presence of obstacles. On the one hand, the offset values to $\xi_{ref}(t)$ directly define the smallest passable cross sections between obstacles. On the other, the effects of noise-related switching between the target points become more and more noticeable as the point density goes down.*

4. Application to the original problem

The proposed trajectory generation is a multi-step process. First, we use the above formulas to generate separate lateral and longitudinal trajectory sets for the different operation modes of the autonomous vehicle. Next, the longitudinal set of each mode will be merged *in every combination* with the lateral set along the center line. The best (with lowest costs) constraint-abiding trajectory can then be identified and executed.

4.1. Generation of the lateral movement

We start with the generation of the lateral movement $d(t)$ as it is more comprehensible than $s(t)$.

When the behavioral layer of the autonomous car considers staying at the lane center (in general the current, for lane changes the adjacent one) the optimal behavior, it simply sets $d_{\text{ref}}(t) \equiv 0$ (see Figure 6, dashed line). With inactive constraints, starting off from a displaced initial state $[d(0), \dot{d}(0), \ddot{d}(0)]^{\mathrm{T}} \neq \mathbf{0}$ and following multiple times the trajectory with the lowest conjoint costs (2) (no matter how the longitudinal movement $s(t)$ and its cost component $J_s(s(t))$ turn out), the trajectory smoothly transfers the vehicle to the reference. In doing so, the proposed optimal policy indeed selects the remainder of the optimal trajectory in the next step, as predicted by Bellman's Principle.

Remark 5. *The more slowly the car drives the more pronounced the non-holonomic constraints (curvature) become in the lateral movement, so that in combination with the longitudinal movement an increasing number of trajectories become invalid. We therefore need to account for the fact that the vehicle can only change its lateral position in combination with a longitudinal movement. More precisely, instead of generating $d(t)$ completely independently of $s(t)$, we switch to $d(s(t))$ below a certain velocity threshold.*

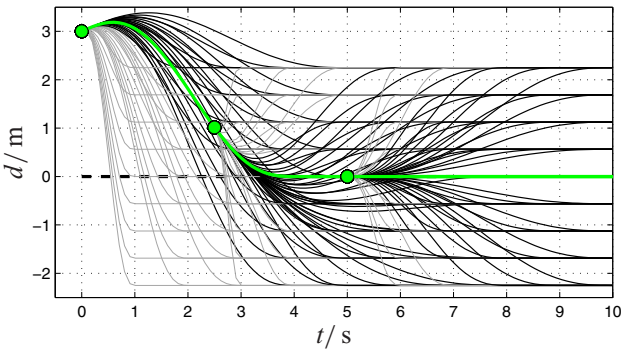

Fig. 6. Simulation of an optimal transfer to the (dashed) reference with $d_{\text{ref}}(t) \equiv 0$ by cyclic replanning. In each step the thick line is the optimal trajectory, with the black lines the valid and the gray lines the invalid alternatives.

Substituting τ by $s(\tau)$, $u = \dddot{d}(t)$ by $(\partial^3/\partial s^3)\, d(s)$ and k_τ by another parameter k_s in (4) *and* (12)*, the quintic polynomial will be retained for an optimal policy for $d(s)$ so that only marginal changes are required in the implementation.*

4.2. Generation of the longitudinal movement

Next, the longitudinal movement is generated, which is required for assuming a certain longitudinal traffic position $s_{\text{ref}}(t)$, such as during stopping at intersections, pulling out into dense traffic on adjacent lanes and following the vehicle ahead with safe distance.

In contrast to the lateral reference with $d_{\text{ref}}(t) \equiv 0$, the longitudinal reference $s_{\text{ref}}(t)$ for merging or following necessitate the prediction (needed in any case for the collision check later on) of the relevant vehicles. For example, knowing the approximate position, velocity and acceleration of the vehicle to follow, we can assume constant acceleration and estimate its future position $s_f(t)$ by integration. The recommended (Schwarzenegger 2010) *constant time gap* then defines the longitudinal reference for our car by

$$s_{\text{ref}}(t) = s_f(t) - [k_0 + k_f\, \dot{s}_f(t)]$$

with suitable parameters $k_0, k_f > 0$. As depicted in Figure 7, we thereby get temporal consistency again, as we subsequently follow the best trajectory tracking the longitudinal reference. For stopping, however, no prediction is required as we simply set $s_{\text{ref}}(t) \equiv$ const. right on the point aimed at.

Other traffic scenarios such as free highways do not require a certain position of the vehicle. In these situations the longitudinal movement is ruled by the recommended vehicle speed. As the lateral velocity component is negligible in practice, we can directly assign it to $\dot{s}_{\text{ref}}(t)$. Due to Theorem 2, it is then sufficient to change over to quartic polynomials and specify only the final speed of the terminal manifold but not the future position, as shown in Figure 8 with $\dot{s}_{\text{ref}}(t) \equiv 5.0\,\text{m/s}$.

For further clarification, an overview of the essential calculations of the simulated operation modes in this section is given in Table 1.

Table 1. Overview of the essential calculations of the simulated operation modes: the uniform diversification of movement in one trajectory set is achieved by various terminal times τ_j and reference offsets δ_i, σ_i and ν_i.

	Target points	Polynomial coefficients	Terminal costs
$d(t)$	$\delta_i \times \tau_j$	$\boldsymbol{q}_{345}(d_{\text{ref}} + \delta_i, 0, 0, \tau_j)$	$k_\tau \tau_j + \frac{1}{2} k_\delta \delta_i^2$
$s(t)$	$\sigma_i \times \tau_j$	$\boldsymbol{q}_{345}([s_{\text{ref}}(\tau_j) + \sigma_i], \dot{s}_{\text{ref}}(\tau_j), \ddot{s}_{\text{ref}}(\tau_j), \tau_j)$	$k_\tau \tau_j + \frac{1}{2} k_\sigma \sigma_i^2$
$\dot{s}(t)$	$\nu_i \times \tau_j$	$\boldsymbol{q}_{34}(\dot{s}_{\text{ref}} + \nu_i, 0, \tau_j)$	$k_\tau \tau_j + \frac{1}{2} k_\nu \nu_i^2$

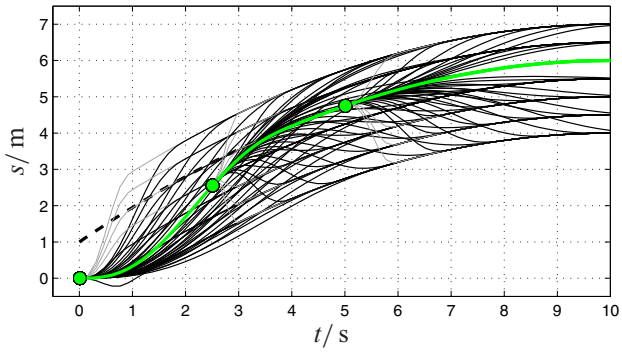

Fig. 7. Simulation of an optimal tracking of the (dashed) longitudinal reference $s_{\text{ref}}(t)$ by cyclic replanning.

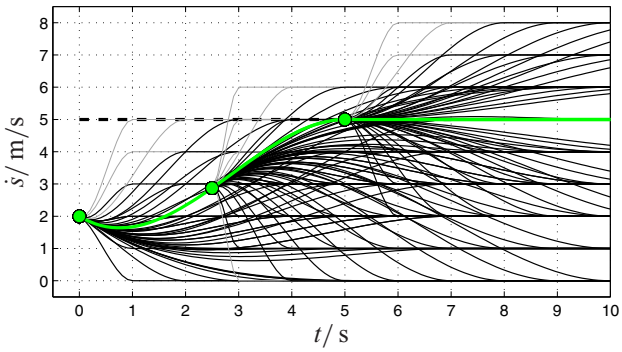

Fig. 8. Simulation of an optimal velocity adaption to the new (dashed) reference velocity $\dot{s}_{\text{ref}}(t) = 5.0\,\text{m/s}$ by cyclic replanning.

4.3. Back transformation to global coordinates

Next, we crosswise superimpose the respective longitudinal with the lateral set and back-transform to global coordinates $\boldsymbol{x}^{\text{T}} = [x_1, x_2]$ (cf. Figure 9), as the adherence to most constraints can only be verified there. With (1) calculating the position $[x_1(t), x_2(t)]^{\text{T}}$ does not pose a problem. However, determining the trajectory's orientation $\theta(t)$, curvature $\kappa(t)$, velocity $v(t)$ and acceleration $a(t)$ (required for the subsequent constraint check and the low-level controller) by consecutive numerical differentiations of (1) is impracticable, as the center line is only approximated by interpolated points in most autonomous architectures. The way out is by utilizing the (also interpolated) center line's

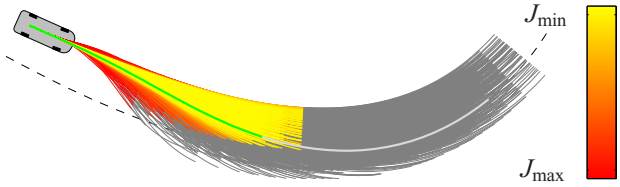

Fig. 9. Superposition of the lateral and longitudinal movements in global coordinates: the trajectory color visualizes (on a horizon of 3.0 s) the respective trajectory costs. As no obstacles are involved, the algorithm gets the vehicle back on the center line and (here) to the desired velocity by following the trajectory with the lowest costs.

orientation, curvature, as well as change of curvature and carrying out the exact closed-form transformation

$$[s, \dot{s}, \ddot{s}; d, \dot{d}, \ddot{d}\,](t) \mapsto [x_1, x_2, \theta, \kappa, v, a](t)$$

in Appendix B.

Remark 6. *Continuous curvatures $\kappa(t)$ and accelerations $a(t)$ of $\boldsymbol{x}(t)$ require a continuous change of curvature of the center line, as can be seen from* (22) *and* (23)*.*

4.4. Fast constraint check

In order to treat all trajectories equally, they have to share the same temporal length. We therefore truncate the long polynomials and extend the short ones on the terminal manifold to the same temporal horizon length (cf. Figure 9). On this horizon we can carry out with the previous closed-form transformations an efficient, pointwise testing of each trajectory for compliance with the maximum admissible curvature and other physically motivated constraints (cf. e. g. Velenis (2006)).

Collision checking with static and dynamic obstacles, however, is much more complex.[5]. In the worst case, *each* trajectory needs to be checked against overlapping with *every* traffic participant and obstacle at *all* times of the considered horizon. An efficient strategy is therefore inevitable.

In order to determine the vehicle's outline relative to each trajectory, we plan at higher speeds for the center of gravity[6] of the vehicle, assuming at this point zero side slip angle. At low speeds, this assumption holds at the rear axis center, so that we switch.

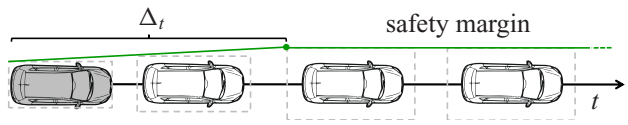

Fig. 10. Gradual increase of safety margin during the first Δ_t of the time horizon providing robustness to measurement noise of the obstacle detection.

With this and a prediction (2nd order) of the moving obstacles, all geometries are approximated by circles allowing for a fast collision check (Ziegler and Stiller 2010).

As a result of the inevitable noise in the obstacle position estimates the 'hard'[7] constraints turn out to be troublesome. Since the trajectory of an obstacle avoidance maneuver deviates only as little from the optimal solution to the free problem as required for maintaining the safety clearance, obstacles are occasionally detected below the safety margin and, at one go, all trajectories get invalidated. As momentary inactivation of the safety margin would only postpone the problem a few planning cycles, we implement another, comparably simple, strategy. As depicted in Figure 10, the safety margin is only enforced after some (fairly short) time Δ_t allowing the trajectories enough time to pull away from every static and moving obstacle.

4.5. Driving mode selection

So far we have implied only one active longitudinal mode. In practice, however, all relevant longitudinal modes are computed in parallel, such as velocity-keeping and stopping. This results for each mode in a single optimal, collision free and feasible trajectory. As the vehicle can only carry out one of them, a suitable selection mechanism is required.

The proposed strategy makes use of the min-operator applied on the signed initial jerk[8] $\dddot{s}(t_0)$ allowing the most retarding (hence most conservative) trajectory to overrule the others (also known as *override control* (Glattfelder and Schaufelberger 2004)). As for approaching a four-way-stop, the velocity mode is active far from the stop position since the optimal position-tracking trajectory would speed up. As the vehicle gets closer, the velocity-keeping trajectory would overrun the crossing so that the position-tracking trajectory takes over and stops.

The selected trajectory is finally carried out by the execution level while the next generation cycle has already started over.

Remark 7. *Even during short cycle times the vehicle considerably changes its position and course at higher speeds. While executing the last step's trajectory, it is therefore required to plan for a future initial state* $[x_1, x_2, \theta, \kappa, v, a](t_0)$. *In general, we rely on the low-level feedback controller so that we use the information from the currently executed trajectory at the future start time* t_0. *In some situations (cf. Figure 18), however, a reinitialization according to the measured movement is required so that the future initial state needs to be calculated from a simple forward simulation of the vehicle's dynamics.*

5. Experiments

The algorithm's ability will now be demonstrated by means of the full-scale experimental vehicle *AnnieWAY* (Kammel et al. 2008) on a tarred runway by means of a simulated environment. The proposed algorithm is implemented in C++ and executed with a cycle time of 200 ms on a single core[9] of the car's main computer (i7 with 3.0 GHz). The generated trajectory is in turn stabilized by multiple low-level feedback controllers (see e. g. Werling et al. (2010a)).

The driving performance (see also Extension 1) is in each of the following eight test scenarios (Figures 11 through 18) visualized by four bird's eye views. In their lower corners the current active longitudinal mode ("v" for velocity and "f" for following/tracking) and the active trajectory number with the total size of the active trajectory set (e. g. "0/4000" means that, according to the cost functional, the best trajectory of 4000 alternatives was selected) are displayed. In addition, this trajectory was tinged according to the color map of Figure 9 on the considered horizon of 3.0 s. Also, the current reference curve in the lane center is plotted as a broken, black line. The current desired position and orientation of the vehicle (the tracking reference of the low-level controller) as well as their predicted values are drawn as rectangular boxes along the current trajectory. The current time is both displayed below the scene and drawn as a broken vertical line in the subjacent signal plots. There, the velocity v, the steering angle δ_s, as well as the tangential and normal vehicle accelerations a_t and a_n are plotted over t.

Simulating the perception and behavioral layer of the autonomous target application, both the predictions of all traffic participants and the reference inputs (target manifolds) are played back throughout the experiment.

5.1. Inner-city test scenario

5.1.1. Driving violation of left-turning traffic At $t = 0$ s the reference curve is placed on the inner lane of the closed traffic loop and the desired speed is set to 8.5 m/s. AnnieWAY therefore accelerates smoothly from its start position along the lane. At the time when the desired speed is reached (Figure 11, top left) the rearmost of the two approaching vehicles on the opposite lane starts unexpectedly an irregular left turn (top right). For collision avoidance, AnnieWAY initially slows down but then takes advantage of the free space between the cars on the other lane (bottom left) so that a full braking is avoided. As soon

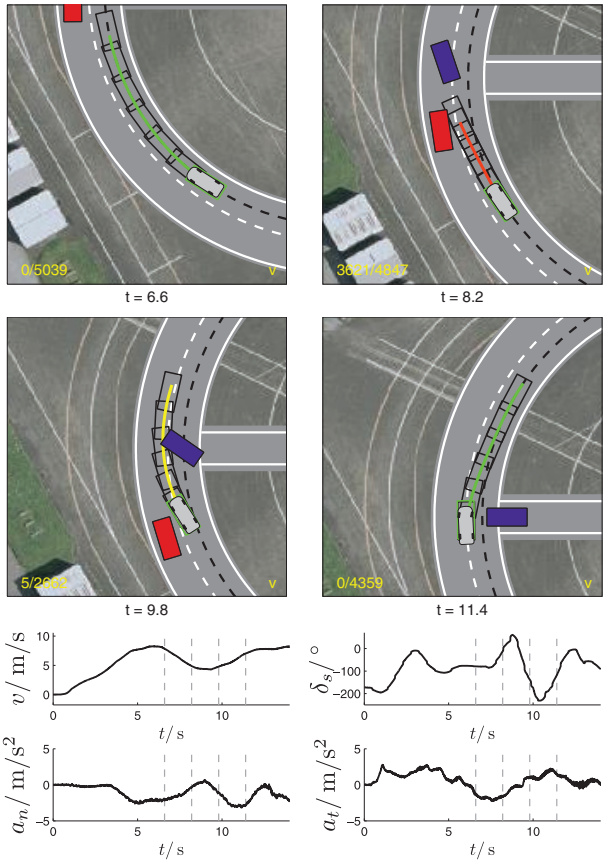

Fig. 11. Driving violation of a left-turning vehicle.

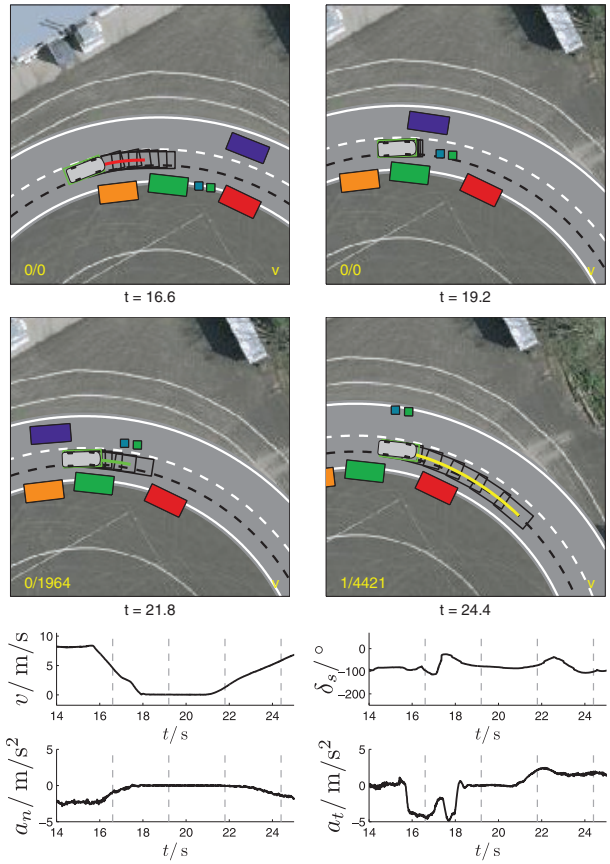

Fig. 12. Pedestrians entering the street between parked cars.

as the turning car has been passed, AnnieWAY returns to the lane center and the desired speed (bottom right).

5.1.2. Pedestrians leaping into view In the subsequent scenario two pedestrians (small squares) enter unforeseen the right road side between parked cars (Figure 12, top left). At this time, AnnieWAY has already deviated from the reference curve in order to maintain the parameterized safety clearance to the parked cars. In contrast to the previous scene, this time the opposite lane is blocked by the approaching vehicle. Additionally, the road crossing pedestrians are 'perceived' fairly late, so that a collision can only be avoided by intensive braking. As none of the trajectories therefore hold the front safety distance, they all become invalid (0/0) and AnnieWAY follows the one to a full stop (top right), which maintains the required clearance the longest. The second the pedestrians clear the street (bottom left), AnnieWAY continues on its way (bottom right).

5.1.3. Side-street traffic offence While AnnieWAY keeps speeding up (Figure 13, top left), another vehicle approaches from the side-street (top right) without slowing down. Avoiding a full braking, AnnieWAY slightly moves sideways, as there is no traffic on the adjacent lane before returning safely to the lane center. Notice, that the autonomous car does not need to significantly deviate from the desired speed (see $v(t)$) during the entire maneuver, as the upcoming vehicle on the opposite lane (bottom right) is still at a safe distance.

5.1.4. Passing of roadside bicyclists Shortly before reaching the initial position after one complete round, two bicyclists (small, oblong rectangles) rush onto the road (Figure 14, top right). Due to the street width AnnieWAY can safely pass them without leaving the lane.

At $t = 36.0\,\mathrm{s}$ the desired speed is set to zero, which puts an end to the autonomous test drive at $t \approx 42.0\,\mathrm{s}$.

5.2. Highway test scenario

At highways new challenges arise, as there the smallest driving mistake can entail a destabilization of the autonomous system with fatal consequences. We therefore put AnnieWAY to the test in a high-speed scenario.

5.2.1. Merging between cars At test start, AnnieWAY travels with $12.5\,\mathrm{m/s}$ in an acceleration lane (Figure 15, top left) and speeds up further as a velocity of $33.3\,\mathrm{m/s}$ is requested from the behavioral layer. In order to pull in into

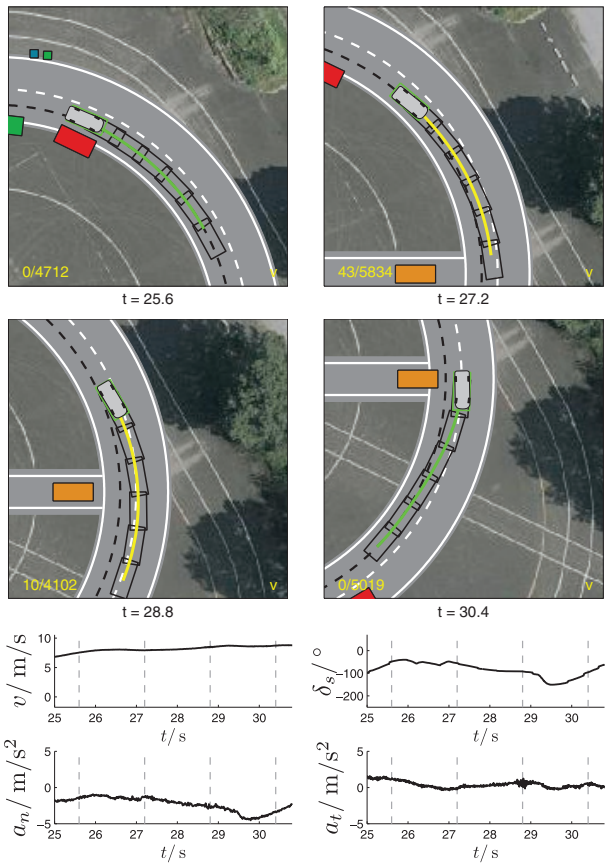

Fig. 13. Driving violation of side-street traffic.

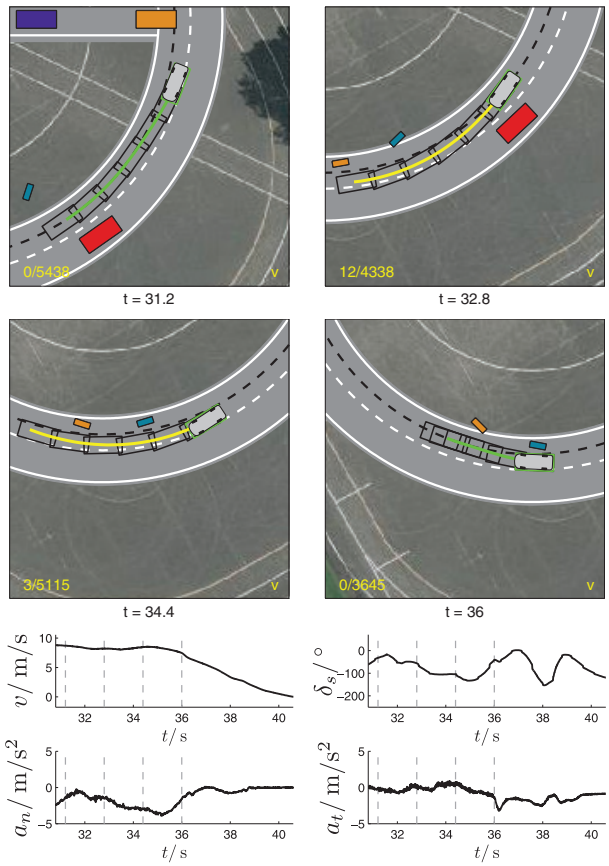

Fig. 14. Passing of bicyclists appearing on the roadside.

the space between the two trucks on the highway (large, oblong rectangles), the longitudinal reference is set right on the projection of this moving spot onto the reference curve. Thereupon AnnieWAY further accelerates while in the velocity mode (v) until the more conservative mode, the position tracking (f), seamlessly takes over (top right) synchronizing the velocity with the gap. At the same time the reference curve switches to the center of the target lane, leading to a very natural merging maneuver (bottom left), which finishes at $t = 20.2$ s (bottom right).

5.2.2. Retarding a lane change The following lane change cannot be smoothly carried out due to a simulated perception error. At $t = 20.4$ s the reference curve shifts by another lane width to the left, initiating the lane change movement (Figure 16, top left). Not until this moment is a quickly approaching car detected on the target lane. In order to maintain the parameterized safety distance[10] the trajectories retard the lane change for another second (top right) so that the maneuver can be completed without serious consequences (bottom left and right).

5.2.3. Veering truck The desired velocity causes the trajectories to further accelerate on the free lane, past the first vehicle (Figure 17, top left), until the truck sideways ahead veers out without prior notice (top right). As the adjacent lane is free, AnnieWAY instantaneously swerves to the left (bottom left) and, as soon as the danger is averted, returns to the reference curve in the center of the middle lane.

5.2.4. Disturbance triggered reinitialization Shortly before reaching the desired speed of 33.3 m/s the course of the vehicle is suddenly altered by a manual steering intervention of a quarter turn, which simulates a swell or a pothole on the road (see δ_s and a_n at $t = 38.5$ m/s in Figure 18). As the low-level tracking controller cannot cope with impulsive disturbances like these, the trajectory simply reinitializes according to the vehicular movement (top right), and the planning algorithm takes advantage of the free adjacent lane. After a few seconds AnnieWAY smoothly returns, similarly to a lane change, to the lane center, where the desired speed is finally reached and the test drive finishes.

6. Discussion

As we have demonstrated, the proposed trajectory generation method copes very well with all of the confronted

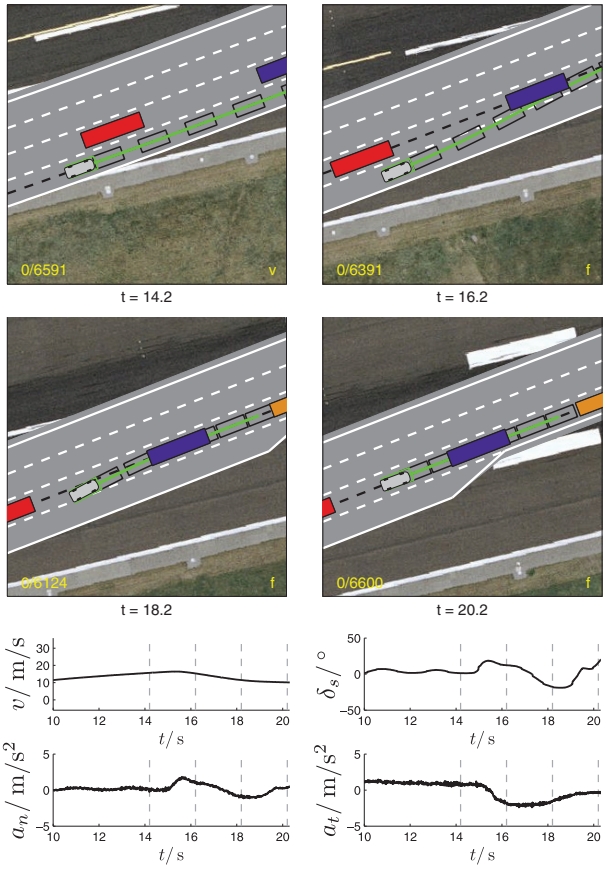

Fig. 15. Merging between cars at high speed.

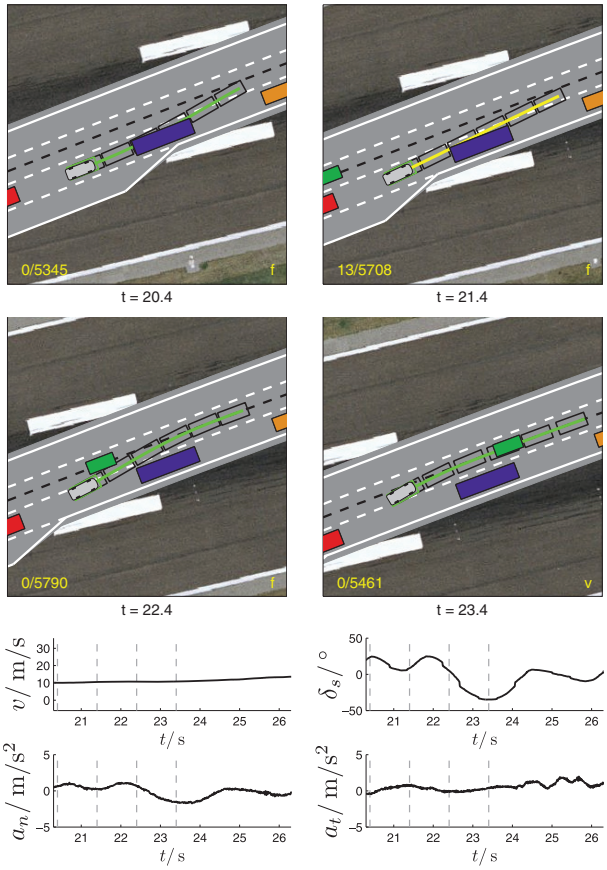

Fig. 16. Collision avoidance by retarding a lane change.

traffic scenarios. While providing for a comfortable, unobtrusive autonomous ride with well-dosed nominal pedal and steering inputs in standard situations like controlled lane changes, stops, and following maneuvers, the algorithm also handles critical, unforeseen situations with aplomb.

Due to the limited dynamics of a vehicle, there are of course situations that inevitably lead to collisions. For understandable reasons, they arise from either extremely late perception (cf. Figure 16), incorrect predictions (cf. Figure 11), and/or the limited optimization horizon (cf. Section 7). As we have surprisingly found out during various simulations, it is much easier to 'corner' the algorithm at low speeds, as the lateral dynamics are then especially limited and braking is only effective for a short time.

Two more aspects are worth mentioning. Despite the wide velocity range of $(0 - 120)$ km/h no parameter adaption is required at all, as the trajectories are velocity invariant. That is, if the vehicular dynamics are not the limiting factor, a velocity adaption of 10 m/s or a lane change of 4 m takes the same time at 30 km/h as at 250 km/h, similar to human driving.

In addition, we referred to the proposed strategy in the beginning as *semi*-reactive. This is because only the constraints relate to a short horizon. The (theoretically infinitely long) target manifold, however, enables the vehicle to initiate a stop sooner and to change lanes slower than the reactive horizon would allow.

7. Conclusion and future work

This paper addresses the online trajectory generation problem faced by an autonomous vehicle in dynamic traffic requiring a combined optimization of the lateral and longitudinal movement. The derived optimal-control-based solution is most suitable for safe lane-changes, distance-keeping, velocity-keeping, merging, etc. amidst moving and stationary obstacles, which we have illustrated by means of a real-time implementation on a test vehicle.

However, the approach is not intended to relieve the behavior layer of making farsighted decisions avoiding inevitable collision states early (see e. g. Martinez-Gomez and Fraichard (2009)). Developing suitable heuristics, such as the well-known *constant-time-gap law*, are essential to minimizing the number of critical situations handled by the reactive layer.

Also, considering uncertainties (see e. g. Althoff and Mergel (2011)) in state estimates, controller performance and, most importantly, traffic prediction has not been touched on by our approach and is left for future research.

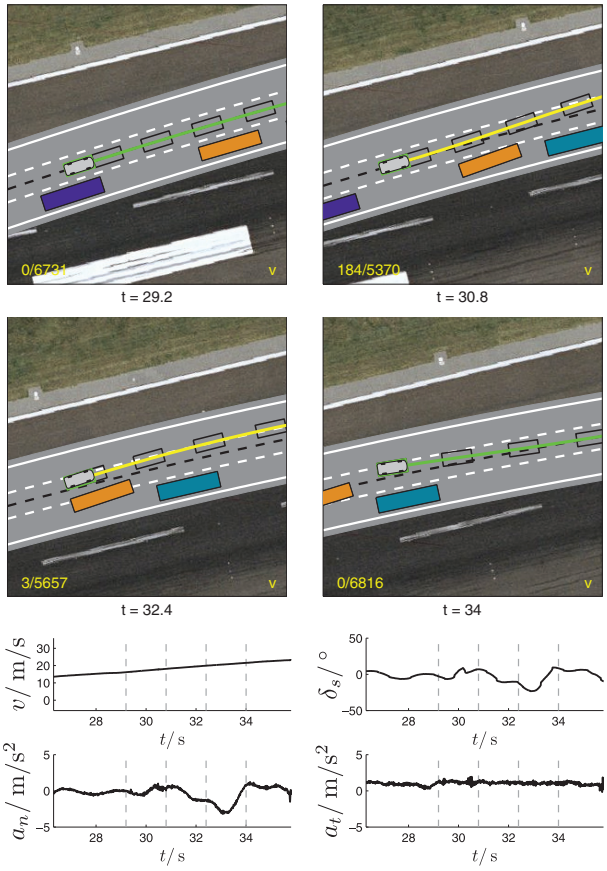

Fig. 17. Overtaking of a veering truck.

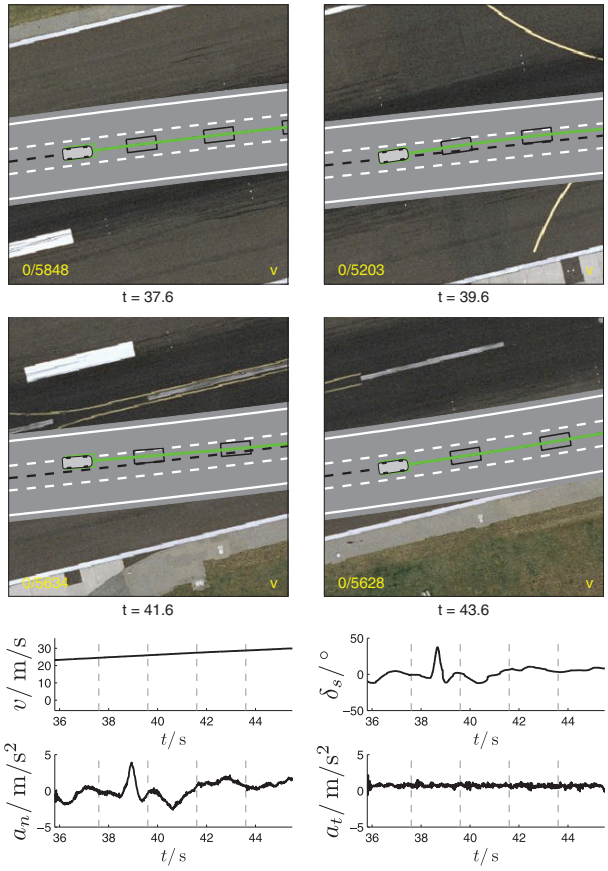

Fig. 18. Disturbance triggered reinitialization after a simulated swell.

Notes

1. *The DARPA Urban Challenge* is a research program conducted in a competitive format to address the challenges of autonomous driving,
   see http://www.darpa.mil/grand challenge.
2. This is a well-known issue of receding-horizon-based methods such as model predictive control (Lee et al. 1998)
3. This becomes clear if one imagines the autonomous car being trapped between four moving vehicles, capable of forcing it onto *any* trajectory by their coordinated motion, e. g. a sine wave.
4. The terminal states would vary in each step if the manifold was discretized relative to the cycle's start time, prohibiting temporal consistency.
5. In our implementation the sole collision checking routine takes in some situations up to 90% of the algorithm's computing time.
6. This offers substantial control advantages exceeding the scope of this paper.
7. In contrast to 'soft' potential field solutions, e. g. Latombe (1990).
8. Due to the proposed generation method, all trajectories match in the lower derivatives, so no distinction can be made there.
9. As the proposed algorithm is entirely parallelizable, the cycle time can be further slashed by multiple cores.
10. Admittedly somewhat small.

Funding

This research was supported by the German Research Foundation (DFG).

Acknowledgments

The authors gratefully acknowledge the cooperation between the German Transregional Collaborative Research Centre 28 *Cognitive Automobiles* and the *ValleyRally* project of Stanford University. Both projects cross-fertilized each other and revealed significant synergy.

References

Althoff M and Mergel A (2011) Comparison of markov chain abstraction and monte carlo simulation for the safety assessment of autonomous cars. *IEEE Transactions on Intelligent Transportation Systems*. To appear.

Bellman R (1954) *The Theory of Dynamic Programming*. Santa Monica, CA: Rand Corporation.

Fletcher L, Teller S, Olson E, Moore D, Kuwata Y, How J, et al. (2008) The MIT-Cornell collision and why it happened. *Journal of Field Robotics* 25(10): 775–807.

Glattfelder A and Schaufelberger W (2004) A path from anti-windup to override control. In *IFAC Symposium on Nonlinear Control Systems*, Stuttgart, Germany, 1379–1384.

Howard T and Kelly A (2007) Optimal rough terrain trajectory generation for wheeled mobile robots. *The International Journal of Robotics Research* 26(2): 141–166.

Kammel S, Ziegler J, Pitzer B, Werling M, Gindele T, Jagszent D, et al. (2008) Team AnnieWAY's autonomous system for the DARPA Urban Challenge 2007. *Journal of Field Robotics* 25(9): 615 – 639.

Kelly A and Nagy B (2003) Reactive nonholonomic trajectory generation via parametric optimal control. *The International Journal of Robotics Research* 22(7–8): 583–601.

Kuwata Y, Fiore GA, Teo J, Frazzoli E and How JP (2008) Motion planning for urban driving using RRT. In *International Conference on Intelligent Robots and Systems*, 1681–1686.

Lacaze A, Moscovitz Y, DeClaris N and Murphy K (1998) Path planning for autonomous vehicles driving over rough terrain. In *Proceedings of the 1998 IEEE ISIC/CIRA/ISAS Joint Conference*, Gaithersburg, MD, September 14-17, 50–55.

Latombe J (1990) *Robot Motion Planning*. Berlin: Springer Verlag.

LaValle M (1998) Rapidly-exploring random trees: A new tool for path planning. Technical report, Computer Science Department, Iowa State University.

LaValle S and Kuffner J (2001) Randomized kinodynamic planning. *The International Journal of Robotics Research* 20(5): 378.

Lee J, Kwon W and Choi J (1998) On stability of constrained receding horizon control with finite terminal weighting matrix. *Automatica* 34(12): 1607 – 1612.

Likhachev M and Ferguson D (2009) Planning long dynamically feasible maneuvers for autonomous vehicles. *The International Journal of Robotics Research* 28(8): 933–945.

Martinez-Gomez L and Fraichard T (2009) Collision avoidance in dynamic environments: an ics-based solution and its comparative evaluation. In: *IEEE International Conference on Robotics and Automation*, 234–241.

Montemerlo M, Becker J, Bhat S, Dahlkamp H, Dolgov D, Ettinger S, et al. (2008) Junior: The Stanford entry in the Urban Challenge. *Journal of Field Robotics* 25(9): 569–597.

Pivtoraiko M and Kelly A (2005) Efficient constrained path planning via search in state lattices. In: *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*.

Schwarzenegger A (2010) *California Driver Handbook*. Department of Motor Vehicles.

Velenis E (2006) *Analysis and control of high-speed wheeled vehicles*. Ph.D. thesis, Georgia Institute of Technology.

Werling M, Gindele T, Jagszent D and Gröll L (2008) A robust algorithm for handling moving traffic in urban scenarios. In: *IEEE Intelligent Vehicles Symposium 2008*, Eindhoven, The Netherlands, 1108–1112.

Werling M, Gröll L and Bretthauer G (2010a) Invariant trajectory tracking with a full-size autonomous road vehicle. *IEEE Transactions on Robotics* 26(4): 758–765.

Werling M, Ziegler J, Kammel S and Thrun S (2010b) Optimal trajectory generation for dynamic street scenarios in a frenet frame. In: *IEEE International Conference on Robotics and Automation*, Anchorage, Alaska, 987–993.

Ziegler J and Stiller C (2010) Fast collision checking for intelligent vehicle motion planning. In: *IEEE Intelligent Vehicles Symposium 2010*, San Diego, USA, 518–522.

A. Index to Multimedia Extensions

The multimedia extension page is found at http://www.ijrr.org

Extension	Media Type	Description
1	Video	Trajectory generation in various situations

B. Closed-form transformations

The biggest challenge of deriving the transformation $[s, \dot{s}, \ddot{s}, d, \dot{d}, \ddot{d}] \mapsto [x_1, x_2, \theta, \kappa, v, a](t)$ is the removal of the singularity in the orientation and curvature at $v = 0$. We tackle this by starting with the transformation $[s, \dot{s}, \ddot{s}, d, d', d''] \mapsto [x_1, x_2, \theta, \kappa, v, a](t)$ with $(\,)' := (\partial/\partial s)$, as it is non-singular at rest and needed at low speeds anyway ($d = d(s)$, see Remark 5).

From Figure 4 we see with $\boldsymbol{R} := [\boldsymbol{t}_c, \boldsymbol{n}_c]$ that

$$\boldsymbol{R}\,[\boldsymbol{x} - \boldsymbol{r}] = [d, 0]^{\mathrm{T}}. \tag{17}$$

With the reference curve's orientation θ_c and curvature $\kappa_c := \mathrm{d}\theta_c/\mathrm{d}s$, the time derivative of (17) then provides $\Delta\theta := \theta - \theta_c$ and the Frenet–Serret formulas the relations

$$\begin{aligned} \dot{s}[1 - \kappa_c d] &= v \cos \Delta\theta \\ \dot{d} &= v \sin \Delta\theta, \end{aligned} \tag{18}$$

so that

$$v = \sqrt{[1 - \kappa_c d]^2 \dot{s}^2 + \dot{d}^2} \tag{19}$$

and

$$d' = \frac{\mathrm{d}}{\mathrm{d}s} d = \frac{\mathrm{d}t}{\mathrm{d}s} \frac{\mathrm{d}}{\mathrm{d}t} d = \frac{\dot{d}}{\dot{s}} = [1 - \kappa_c d] \tan \Delta\theta \tag{20}$$

holds.

With the covered arc length s_x of the trajectory we get the relation

$$\frac{\mathrm{d}}{\mathrm{d}s} = \frac{\mathrm{d}s_x}{\mathrm{d}s} \frac{\mathrm{d}}{\mathrm{d}s_x} = \frac{\mathrm{d}s_x}{\mathrm{d}t} \frac{\mathrm{d}t}{\mathrm{d}s} \frac{\mathrm{d}}{\mathrm{d}s_x} = \frac{v}{\dot{s}} \frac{\mathrm{d}}{\mathrm{d}s_x} = \frac{1 - \kappa_c d}{\cos \Delta\theta} \frac{\mathrm{d}}{\mathrm{d}s_x}.$$

With $\kappa := \frac{\mathrm{d}}{\mathrm{d}s_x}\theta$ deriving $\Delta\theta$ for s therefore yields

$$\frac{\mathrm{d}}{\mathrm{d}s} \Delta\theta = \left[\kappa \frac{1 - \kappa_c d}{\cos \Delta\theta} - \kappa_c \right]. \tag{21}$$

Deriving (20) once more, we therefore get

$$d'' = -[\kappa_c' d + \kappa_c d'] \tan \Delta\theta + \frac{1 - \kappa_c d}{\cos^2 \Delta\theta} \left[\kappa \frac{1 - \kappa_c d}{\cos \Delta\theta} - \kappa_c \right]. \tag{22}$$

As the vehicle travels along the road, we can assume that $|\Delta\theta| < \frac{\pi}{2}$ and $1-\kappa_c d > 0$, and equations (20) and (22) can be solved for θ and κ, including the special case $v = 0$. Another time derivative finally provides with (21) the longitudinal acceleration

$$a := \dot{v} = \ddot{s}\frac{1-\kappa_c d}{\cos\Delta\theta} + \dot{s}\frac{\mathrm{d}s}{\mathrm{d}t}\frac{\mathrm{d}}{\mathrm{d}s}\frac{1-\kappa_c d}{\cos\Delta\theta} = \ddot{s}\frac{1-\kappa_c d}{\cos\Delta\theta} + \frac{\dot{s}^2}{\cos\Delta\theta}\left[[1-\kappa_c d]\tan\Delta\theta\left[\kappa\frac{1-\kappa_c d}{\cos\Delta\theta} - \kappa_c\right] - [\kappa_c' d + \kappa_c d']\right]. \tag{23}$$

As for the transformation for $d(t)$ at higher speeds where $\dot{s} \neq 0$, the relations

$$\dot{d} = \frac{\mathrm{d}}{\mathrm{d}t}d = \frac{\mathrm{d}s}{\mathrm{d}t}\frac{\mathrm{d}}{\mathrm{d}s}d = \dot{s}d', \tag{24}$$

$$\ddot{d} = \frac{\mathrm{d}}{\mathrm{d}t}\dot{s}d' = \ddot{s}d' + \dot{s}\frac{\mathrm{d}s}{\mathrm{d}t}\frac{\mathrm{d}}{\mathrm{d}s}d' = \ddot{s}d' + \dot{s}^2 d'' \tag{25}$$

can be sequentially solved for d' and d'' and substituted in the previous transformations. □