

Modified Dynamic Programming Algorithms for GLOSA Systems with Stochastic Signal Switching Times

Panagiotis Typaldos^{a,*}, Markos Papageorgiou^{a,b}

^a*Dynamic Systems and Simulation Laboratory, School of Production Engineering & Management, Technical University of Crete, 73100, Chania, Greece*

^b*Faculty of Maritime and Transportation, Ningbo University, Ningbo, China*

ARTICLE INFO

Keywords:

GLOSA

Stochastic Dynamic Programming

Discrete Differential Dynamic Programming (DDDP)

Differential Dynamic Programming (DDP)

Traffic Light Advisory

Speed Advisory

ABSTRACT

A discrete-time stochastic optimal control problem was recently proposed to address the GLOSA (Green Light Optimal Speed Advisory) problem in cases where the next signal switching time is decided in real time and is therefore uncertain in advance. The corresponding numerical solution via SDP (Stochastic Dynamic Programming) calls for substantial computation time, which excludes problem solution in the vehicle's on-board computer in real time. In this context, the present paper concentrates on the challenge of developing numerical algorithms to solve efficiently the stochastic GLOSA problem. As a first attempt to overcome the computation time bottleneck, a modified version of Dynamic Programming, known as Discrete Differential Dynamic Programming (DDDP) was recently employed for the numerical solution of the stochastic optimal control problem and was demonstrated to achieve results equivalent to those obtained with the ordinary SDP algorithm, albeit with significantly reduced computation times. After an outline of the DDDP approach, the present work considers a second modified version of Dynamic Programming, known as Differential Dynamic Programming (DDP). For the stochastic GLOSA problem, it is demonstrated that DDP achieves quasi-instantaneous (extremely fast) solutions in terms of CPU times, which allows for the proposed approach to be readily executable online, in an MPC (Model Predictive Control) framework, in the vehicle's on-board computer. The novel numerical approach is tested and demonstrated by use of realistic examples and is compared to the SDP and DDDP solutions. It should be noted that DDP does not require discretization of variables, hence the obtained solutions may be slightly superior to the standard SDP solutions.

1. Introduction

In an era of climate crisis, actions toward environment protection are due in Transportation systems, which have a significant impact on environmental pollution and energy resources shortage. Improved fuel efficiency in road vehicles implies reduced fuel consumption, emissions and expenses for the driver. Efficient intelligent transportation systems, including real-time traffic signals, contribute to this effort. Traffic lights were introduced, in the first place, as a measure to ensure safe crossing of antagonistic streams of vehicles, cyclists and pedestrians at urban intersections. However, enforcing safety via traffic lights affects the fuel consumption, as some vehicles are obliged to fully stop in front of a red light; and then accelerate, after the traffic light switching to green. To reduce the resulting vehicle delays and number of stops, several algorithms have been proposed and deployed over the past decades, aiming at optimizing the traffic signals operation (Hounsell and McDonald, 2001; Papageorgiou et al., 2003; Kosmatopoulos et al., 2006). In this context, fuel consumption is increasingly considered as an optimization or evaluation criterion, while developing and deploying signal control systems (Jamshidnejad et al., 2017).

Fixed-time signal plans are derived offline based on historical constant demands; and are applied without deviations. This implies that switching times of the traffic lights are known in advance. In contrast, real-time signal control strategies make use of real-time measurements to compute in real time suitable signal settings. Depending on the employed signal control strategy, the control update period may range from one second to one signal cycle. Clearly, for real-time signals, the next switching time is not known before the switching decision has been actually made.

A common dilemma, for a vehicle approaching a red traffic light, is whether it should slow down or maintain its speed at risk to stop if the traffic light has not switched to green at arrival. This dilemma may be addressed by appropriately designed on-board systems, which receive the current state and timing of the traffic signal and guide the

*Corresponding author
ORCID(s):

driver (or an automated vehicle), while approaching the traffic light, by giving speed advice that ensures that the vehicle will cross the traffic signal at green and with minimum fuel consumption and emissions. Systems (or apps) optimizing the vehicle approach to traffic lights are often referred to as Green Light Optimal Speed Advisory (GLOSA) systems (Stahlmann et al., 2016). GLOSA apps work on the assumption that there is no slower vehicle in front, otherwise the GLOSA advice cannot be followed.

In the case of fixed signals, and hence prior knowledge of the next switching time, a corresponding message is broadcasted by the signal controller well before switching. Under these conditions, the problem of how to optimize the approach to traffic signals has been addressed in different ways, including rule-based algorithms (Katsaros et al., 2011; Sanchez et al., 2006; Ma et al., 2018) and optimal control approaches, which consider explicitly the vehicle kinematics and are, by their nature, more efficient in producing fuel-optimal speed profiles (Lawitzky et al., 2013; Typaldos et al., 2020a). Some of these works take into consideration additional factors, such as the car waiting queue. In (Yang et al., 2016), an eco-CACC (Cooperative Adaptive Cruise Control) scheme was developed, where optimal speed trajectories were calculated by minimizing fuel-consumption and ensuring that the vehicles arrive at the intersection as soon as the last queued vehicle is discharged. An extension of this work for the case of multiple signalized intersections is presented in (Yang et al., 2020). In the literature, most existing GLOSA problems are focusing on developing strategies for a single vehicle, however, there are few works that approach the problem in case of platoons of vehicles in mixed traffic environment. In (Zhao et al., 2018), a model predictive control (MPC) approach was proposed that minimizes the fuel-consumption of a platoon of mixed vehicles (automated and human-driven vehicles) while crossing an intersection during the green phase.

On the other hand, the case of real-time signals with very short (e.g., second-by-second) control update periods is more complicated, as there is no exact prior knowledge of the next switching time. In this case, the best available knowledge can be presented as an estimate or as a probabilistic distribution for the next switching time within a short-term future time-window. A comparison of several signal control methods with respect to traffic efficiency, plan stability and the resulting GLOSA performance is presented in (Blokpoel and Lu, 2017). In (Koukoumidis et al., 2011), a system called SignalGuru was developed, which detects and predicts the traffic signal timing using a collection of mobile phone data. Bodenheimer et al. (2014) developed a methodology based on graph theory to forecast fully and semi-adaptive traffic signal changes. Similarly, Eteifa et al. (2021) and Scheegans et al. (2022) used Artificial Intelligence (AI) and learning techniques in order to estimate the switching times of actuated traffic signals. Instead of a single estimate, Mahler and Vahidi (2012) produced a probability distribution for the next switching time based on past signal statistics. However, the available probability distribution is used heuristically (rather than optimally) to accordingly time-weight the objective function within a deterministic optimal control problem that is solved via a dynamic programming (DP) algorithm. Hao et al. (2018) developed an Eco-Approach and Departure (EAD) approach, that is compatible with actuated signals. The traffic signal sends to the vehicle a minimum and maximum time-to-change of the respective signal phase and, based on this window, the method recommends an appropriate speed via a rule-based procedure. In (Sun et al., 2020), the uncertainty of traffic signal timing is considered as a data-driven chance-constrained robust optimization problem, where the red-light duration is defined as a random variable. A data-driven approach is adopted to formulate chance constraints, based on empirical sample data, and Dynamic Programming is employed to solve the optimization problem. In (Lawitzky et al., 2013), a stochastic optimal control problem was proposed, which uses a probabilistic distribution of the next signal switching time, and the problem is solved using a discrete stochastic dynamic programming (SDP) algorithm.

In summary, there are several GLOSA apps in various countries for the case of fixed signals, but in case of real-time adaptive signals, where the next switching time is not known, existing apps are structurally not applicable. Motivated by this lack, we developed, in a previous work (Typaldos et al., 2020a), a novel Stochastic Dynamic Programming (SDP) approach for the GLOSA problem, for the case of traffic signals with unknown switching times. Specifically, in (Typaldos et al., 2020a), the problem of producing fuel-optimal vehicle trajectories for a vehicle approaching a traffic signal was considered for both cases of known and stochastic switching times. For the first case, an optimal control problem was formulated and solved analytically via PMP (Pontryagin's Maximum Principle). Subsequently, the case of stochastic switching time with known probability distribution was also addressed in the format of a stochastic optimal control problem, which was solved numerically using SDP. The proposed SDP algorithm may take several minutes to execute, which implies that the solution is not real-time feasible and can therefore not be obtained on-board the vehicle, but must be executed offline, on the infrastructure side, and be communicated to approaching vehicles according to their current states. To substantially reduce the computation time and memory requirements for the solution of the stochastic GLOSA problem and enable its solution on-board the vehicle, Typaldos et al. (2021) employed a Discrete

Differential Dynamic Programming (DDDP) algorithm (Heidari et al., 1971). DDDP is an iterative algorithm, each iteration receiving a feasible (but non-optimal) trajectory and transforming it to an enhanced one, to be used in the next iteration, and so forth, until convergence to the overall solution is achieved. Each iteration uses the conventional DP algorithm to solve the problem within a strongly reduced state domain around the received state trajectory of the previous iteration. The computational time required for each iteration is far lower compared to that of the one-shot full problem solution, due to the strongly reduced space domain considered. Thus, the problem may be solved by use of DDDP very fast (around 1 s CPU time or less) (Typaldos et al., 2021). It should be noted that, in contrast to deterministic optimal control problems, general stochastic optimal control problems do not feature a solution trajectory, even for given initial states, due to the uncertainty of system evolution created by the included stochastic variables. However, in the specific stochastic GLOSA problem, such a trajectory is indeed present in the problem solution, and this allows for application of the DDDP algorithm.

As an attempt to further reduce the computation time for the solution of the stochastic GLOSA problem and enable its solution on-board the vehicle, a Differential Dynamic Programming (DDP) algorithm is employed in the current work. DDP was first introduced in (Jacobson and Mayne, 1970), while several extensions and applications were eventually conducted by Murray and Yakowitz (1979, 1984). DDP is also used in robotics (Tassa et al., 2014; Budhiraja et al., 2018) and unmanned aerial vehicles (UAV) (Xie et al., 2017). DDP is also an iterative algorithm, which starts with a user-defined feasible trajectory to be used in the first iteration. Each iteration produces an improved trajectory, to be used in the subsequent iteration; until convergence to the solution of the original optimal control problem is achieved. Each iteration solves, for each time-step separately, a quadratic-linear approximation of the stochastic recursive Bellman equation; the approximation being computed around the initial trajectory of the iteration. It should be noted that, in contrast to the standard DP and DDDP algorithms, DDP requires no discretization of the state and control spaces, hence it may result in better solutions compared with discretised-value formulations. Moreover, DDP is found to deliver solutions quasi-instantaneously, which enables the approach to be readily implementable in an MPC framework online, on the vehicle's computer. Note that communication between vehicles and traffic signals (V2I) is crucial in GLOSA problems (Lu et al., 2020), however in this work we focus on the methodological part of such systems. In other words, by its very nature, the present paper concentrates on challenging algorithmic issues, specifically on how to solve efficiently the stochastic GLOSA problem.

The remainder of the paper is organized as follows: in Section 2, the optimal control problems with known signal switching time and uncertain signal switching time, as proposed in (Typaldos et al., 2020a) are briefly presented for completeness. Section 3 presents the DDDP algorithm, along the lines of (Typaldos et al., 2021). Section 4 presents the DDP-based GLOSA problem algorithm, while demonstration results for the DDP algorithm's performance and comparison with the one-shot stochastic GLOSA and the DDDP results are presented in Section 5. Section 6 discusses the features, limitations and usage of the work. Finally, Section 7 concludes this work, summarizing its contributions.

2. Optimal Control-Based GLOSA with Known or Unknown Signal Switching Time

This section outlines, for completeness, the deterministic and stochastic GLOSA optimal control problems, for the cases of known and unknown signal switching time, respectively, as well as the SDP solution algorithm for the latter, following (Typaldos et al., 2020a).

2.1. Known Signal Switching Time - Deterministic GLOSA

Consider a vehicle traveling from an initial state $\mathbf{x}_0 = [x_0, v_0]^T$, x_0 being a given initial position and v_0 a given initial speed, with the purpose to reach a fixed final state $\mathbf{x}_e = [x_e, v_e]^T$ within a free (but weighted) time horizon t_e . Between the initial and final positions, there is a traffic signal, and hence the additional restriction that the vehicle must not pass the signal position x_1 before the known time t_1 , when the traffic light turns green from red (see Fig. 1). The implicit assumption here is that a red light is active when the vehicle appears on the link (at time 0), but a generalisation, which includes the case where the vehicle appears when the traffic light is green, is given in (Typaldos et al., 2020a). The vehicle task is to appropriately adjust its acceleration (control variable), so as to minimize fuel consumption, while satisfying the initial and final conditions \mathbf{x}_0 and \mathbf{x}_e , as well as the intermediate (traffic signal) constraint.

The minimization problem outlined above is formulated as an optimal control problem, which accounts for the vehicle kinematics via the following state equations:

$$\dot{\mathbf{x}} = \mathbf{v} \quad (1)$$




Figure 1: A vehicle starting from a given initial state (x_0, v_0) aims to reach a fixed final state (x_e, v_e) in presence of a traffic signal at a given position x_1 that switches to green at time t_1 .

$$\dot{v} = a \quad (2)$$

where a is the vehicle acceleration, which assumes the role of the control variable. The objective criterion to be minimized reads

$$J = w \cdot t_e + \frac{1}{2} \int_0^{t_e} a^2 dt \quad (3)$$

The utilized acceleration cost term a^2 in the criterion was demonstrated to be an excellent proxy for deriving fuel-minimizing vehicle trajectories (Typaldos et al., 2020b). The final time t_e is free but penalized with the weight w . For higher values of w , the resulting t_e will be lower and vice versa. This, consequently, affects the acceleration cost, which, depending on higher or lower w value, will also have increased or decreased values (for more details, see (Typaldos et al., 2020a)). In addition, the green-light constraint, $t_S \geq t_1$, must be fulfilled, where t_S is the time at which the vehicle crosses the signal position x_1 , that is, $x(t_S) = x_1$. If necessary, upper and lower bounds may be applied to speed v and acceleration a .

The solution of this problem was addressed in (Typaldos et al., 2020a) and can be obtained analytically using symbolic differentiation tools. Thus, the numerical solution of the deterministic GLOSA problem, for a specific problem instance, takes only fractions of a second of computation time and can be executed in real time on-board for each approaching vehicle. Note that it may be reasonable to continuously update the vehicle trajectory (in an MPC loop) to account for possible deviations from the first calculated vehicle trajectory, which may occur because of a variety of disturbances, including a slower vehicle ahead.

For a given intersection, the final state is the same for any initial vehicle state \mathbf{x}_0 and any switching time t_1 . Therefore, the optimal value of criterion (3) of the deterministic GLOSA solution depends only on the initial state and the switching time and is denoted $J_{DG}^*(\mathbf{x}_0, t_1)$ for later use.

2.2. Uncertain Signal Switching Time Problem – Stochastic GLOSA

The traffic light switching time may be subject to short-term decisions in dependence of the prevailing traffic conditions in cases of real-time signals. In such cases, we typically have minimum and maximum admissible switching times; hence, based on statistics from past signal switching activity, we may derive a probability distribution of switching times within the admissible time-window of possible signal switching times. Thus, the problem can be cast in the format of a stochastic optimal control problem, which may be solved numerically using SDP techniques. To this end, the analytical solution of the deterministic GLOSA optimal control problem is used within the stochastic approach, as will be explained in this section (for more details, see (Typaldos et al., 2020a)).

For the stochastic optimal control problem and the SDP algorithm (Bertsekas, 2012), the discrete-time version of the vehicle kinematics, with time step T , is considered, as follows:

$$x(k+1) = x(k) + v(k)T + \frac{1}{2}a(k)T^2 \quad (4)$$

$$v(k+1) = v(k) + a(k)T \quad (5)$$

where $x(k)$ and $v(k)$ correspond to the vehicle position and speed, respectively, at discrete times $k = 0, 1, \dots$ (where $kT = t$), while the control variable $a(k)$ reflects the acceleration that remains constant over each time period k . The state and control variables are bounded within the following feasible regions

$$\mathbf{x}(k) \in \mathbf{X} = [\mathbf{x}_{min}, \mathbf{x}_{max}] \quad (6)$$

$$a(k) \in U = [a_{\min}, a_{\max}] \quad (7)$$

with \mathbf{x}_{\min} , \mathbf{x}_{\max} and a_{\min} , a_{\max} being the lower and upper bounds of the states and acceleration, respectively. The traffic light's discrete switching time k_1 is not known, but it is assumed that a known range $k_{\min} \leq k_1 \leq k_{\max}$ of possible switching times exists, k_{\min} and k_{\max} being the minimum and maximum possible switching times, respectively.

For proper problem formulation, a virtual state $\tilde{x}(k)$ is introduced, that reflects formally the stochasticity of traffic light switching

$$\begin{aligned} \tilde{x}(k+1) &= \tilde{x}(k) \cdot z(k) \\ \tilde{x}(0) &= 1 \end{aligned} \quad (8)$$

where $z(k)$ is a binary stochastic variable defined as

$$z(k) = \begin{cases} 0 & \text{if traffic light switches at time } k+1 \\ 1 & \text{else} \end{cases} \quad (9)$$

with (8) and (9), the virtual state $\tilde{x}(k)$ is either equal to 1 if the traffic light has not yet switched until time $k-1$; or equal to 0 if switching occurred at time k or earlier. The virtual state \tilde{x} is assumed measurable, which means that the system knows at each time kT if switching has taken place or not within the last time period $[(k-1)T, kT]$.

The stochastic variable $z(k)$ is independent of its previous values and takes values according to a time-dependent probability distribution $p(z|k)$. Based on the statistics of previous signal switching activity, availability of an a-priori discrete probability distribution $P(k)$, $k_{\min} \leq k_1 \leq k_{\max}$, is assumed, for signal switching within the time-window, where $\sum_{k=k_{\min}}^{k_{\max}} P(k) = 1$. Since no switching takes place for $k \leq k_{\min} - 1$, we have

$$p(0|k) = 0 \quad \text{for } k < k_{\min} - 1 \quad (10)$$

For $k \geq k_{\min}$, the probability distribution $p(z|k)$, is obtained by use of "crop-and-scale", meaning that the a-priori probabilities of previous time steps, where switching did not take place, are distributed analogously to increase the probabilities of the remaining discrete times within the time-window (Lawitzky et al., 2013). As shown in (Typaldos et al., 2020a), this update may be done by use of the following crop-and-scale formula that applies for $k_{\min} \leq k \leq k_{\max} - 1$ and for any a-priori distribution $P(k)$

$$p(0|k) = \frac{P(k+1)}{\sum_{\kappa=k+1}^{k_{\max}} P(\kappa)} \quad (11)$$

The cost criterion of the stochastic problem is essentially the same as (3) in the deterministic case. However, in the stochastic case, the exact value of the criterion depends on the stochastic variable's realization, and therefore we consider minimisation of the expected value. As shown in (Typaldos et al., 2020a), the cost criterion in discrete time becomes

$$J = E \left\{ \frac{1}{2} \sum_{k=0}^{k_1-1} a(k)^2 + J_{DG}^*[\mathbf{x}(k_1), k_1] \right\} \quad (12)$$

where the expectation refers to the stochastic variable $z(k)$, $k = 0, \dots, k_{\max} - 1$. This criterion captures the quadratic acceleration cost up until the stochastic switching time k_1 , when the vehicle is at state $\mathbf{x}(k_1)$. After the signal switching, the problem instantly becomes a deterministic GLOSA problem (Section 2.1), and the corresponding optimal cost-to-go is $J_{DG}^*[\mathbf{x}(k_1), k_1]$, which will be denoted as the "escape cost".

To obtain a formally proper cost criterion, the stochastic variable $z(k)$ and the virtual variable $\tilde{x}(k)$, introduced earlier, are used, and, as shown in (Typaldos et al., 2020a), this yields the objective function in the required form, as follows

$$J = E \left\{ \tilde{x}(k) \sum_{k=0}^{k_{\max}-1} \left[\frac{1}{2} a(k)^2 + [1 - z(k)] J_{DG}^*[\mathbf{x}(k), a(k), k+1] \right] \right\} \quad (13)$$

where, for notational simplicity, $J_{DG}^*[\mathbf{x}(k), a(k), k+1] = J_{DG}^*[\mathbf{x}(k+1), k+1]$ when $\mathbf{x}(k+1)$ is replaced via the state equations (4), (5) as a function of $\mathbf{x}(k), a(k)$.

Equations (4)-(13) constitute an ordinary stochastic optimal control problem (Bertsekas, 2012). Denoting the corresponding optimal cost-to-go function by $V[\mathbf{x}(k), \tilde{x}(k), k]$, the recursive Bellman equation for $0 \leq k \leq k_{max} - 1$ reads

$$\begin{aligned} V[\mathbf{x}(k), \tilde{x}(k), k] = \min_{a(k) \in U} \left\{ E \left\{ \frac{1}{2} a(k)^2 + [1 - z(k)] J_{DG}^*[\mathbf{x}(k), a(k), k+1] \right. \right. \\ \left. \left. + V[\mathbf{x}(k+1), \tilde{x}(k)z(k), k+1] \right\} \right\} = \min_{a(k) \in U} \left\{ \frac{1}{2} a(k)^2 + p(0|k) \cdot J_{DG}^*[\mathbf{x}(k), a(k), k+1] \right. \\ \left. + [1 - p(0|k)] \cdot V[\mathbf{x}(k+1), 1, k+1] \right\} \end{aligned} \quad (14)$$

with boundary condition $V[\mathbf{x}(k_{max}), 1, k_{max}] = 0$. Note that, the minimum is required with respect to the control of time k only, i.e., $a(k) \in U$, as typical in dynamic programming, which facilitates the numerical solution.

In this formulation, it has been assumed that the decision on the traffic light switching is taken within the last time period before the actual switching. The generalization to the case of taking a switching decision κ time periods in advance of the actual switching is easy and is mentioned in (Typaldos et al., 2020a).

2.3. Discrete SDP Numerical Solution Algorithm

For the numerical solution of the stochastic problem, using the SDP algorithm, the state and control variables must be discretised. As the discretization level has a significant impact on computational time and memory requirements, but also on the accuracy of the computed solution, an appropriate trade-off must be specified between reasonable computation requirements versus achievable solution quality. For the discretization, the discrete time interval T is set equal to 1 s, which is a reasonable choice for the problem at hand. Then, a general discretization interval Δ for the problem variables is assumed, and the discretization interval of acceleration is set $\Delta a = \Delta$. From (5), the discretization interval of speed assumes the same value ($\Delta v = \Delta a = \Delta$). Likewise, in view of (4), the discretization interval for the position is

$$\Delta x = \frac{1}{2} \Delta \cdot T^2 \quad (15)$$

Based on these settings, it was shown in (Typaldos et al., 2020a) that, if $x(k), v(k), a(k)$ are discrete points, then $x(k+1)$ and $v(k+1)$ resulting from (4) and (5) are also discrete points.

It is now straightforward to apply the discrete SDP algorithm to obtain a globally optimal closed-loop control law $a^*(k) = R[\mathbf{x}(k), k]$, which, for any given vehicle state $\mathbf{x}(k) \in \mathbf{X}$ and time k , delivers the corresponding optimal acceleration $a^*(k)$. The SDP algorithmic steps are summarized below. Note that the algorithm needs to consider only the case $\tilde{x}(k) = 1$, therefore any arguments pertaining to $\tilde{x}(k)$ are suppressed for convenience.

The SDP algorithm is as follows:

- 1: $V[\mathbf{x}(k_{max}), k_{max}] \leftarrow 0 \quad \forall \mathbf{x}(k_{max}) \in \mathbf{X}$
- 2: **for** each $k = k_{max} - 1, \dots, 0$ **do**
- 3: **for** each discrete state $\mathbf{x}(k) \in \mathbf{X}$ **do**
- 4: **for** each discrete control $a(k) \in U$ **do**
- 5: Calculate $x(k+1), v(k+1)$
- 6: **if** $\mathbf{x}(k+1) \notin \mathbf{X}$ **then**
- 7: $\Phi[\mathbf{x}(k), a(k), k+1] \leftarrow \infty$
- 8: **continue**
- 9: **end if**
- 10: $\Phi[\mathbf{x}(k), a(k), k+1] \leftarrow \frac{1}{2} a(k)^2 + p(0|k) \cdot J_{DG}^*[\mathbf{x}(k), a(k), k+1]$
 $\quad + [1 - p(0|k)] \cdot V[\mathbf{x}(k+1), k+1]$
- 11: **end for**
- 12: $V[\mathbf{x}(k), k] \leftarrow \min \Phi[\mathbf{x}(k), a(k), k+1] \quad \forall a(k) \in U$
- 13: $R[\mathbf{x}(k), k] = a^*(k) \leftarrow \arg \min_{a(k) \in U} \{ \Phi[\mathbf{x}(k), a(k), k+1] \},$

```

14:         with  $a^*(k)$  the optimal control of point  $[\mathbf{x}(k), k]$ 
15:     end for
16: end for
    
```

As mentioned, this algorithm delivers a globally optimal control law $a^*(k) = R[\mathbf{x}(k), k]$ for the full state domain \mathbf{X} . Note that general stochastic optimal control problems do not possess a solution trajectory for specific initial states \mathbf{x}_0 , because the state evolution is uncertain in presence of the stochastic variables. However, in the specific stochastic GLOSA problem addressed here, the evolution of the vehicle states (4) and (5) is not affected by the stochastic variable $z(k)$, which concerns only the signal switching time. Thus, for a given initial state, i.e., vehicle position and speed at time 0, the optimal control law may be used to produce an optimal vehicle trajectory that the vehicle should pursue; until the signal switching actually occurs, in which case the vehicle must instantly behave according to the deterministic GLOSA solution.

Note that the SDP algorithm may take several minutes to execute in an ordinary PC, as will be reported in Section 5. This implies that the solution cannot be obtained on the vehicle side in real time.

3. Discrete Differential Dynamic Programming (DDDP)

In (Typaldos et al., 2021), a modified DP algorithm called DDDP was employed for the stochastic GLOSA problem in order to address the major disadvantage of the discrete (S)DP algorithm, which is the high computation time required for the numerical solution of the optimal control problem, which increases exponentially with the number of state and control variables (Papageorgiou et al., 2015). The DDDP algorithm was proposed in (Heidari et al., 1971) for deterministic optimal control problems. The method can nevertheless be applied here, because an optimal vehicle trajectory may be derived for the stochastic GLOSA problem, despite its stochastic character, as mentioned in the previous section. The DDDP approach to the stochastic GLOSA problem is presented in this paper in more detail and with more comprehensive testing and comparison results, compared to the conference article (Typaldos et al., 2021).

DDDP is an iterative algorithm, calling for a feasible starting state trajectory to be specified externally. Each iteration l receives a feasible (but non-optimal) state trajectory $\mathbf{x}^{(l-1)}(k)$, and transforms it to an enhanced one $\mathbf{x}^{(l)}(k)$, to be used in the next iteration. To this end, each iteration solves a discrete SDP problem by use of the standard SDP algorithm presented above. What changes at each iteration l is the considered state domain

$$\mathbf{X}_C^{(l)} = \left\{ \mathbf{x}(k) \mid |\mathbf{x}(k) - \mathbf{x}^{(l-1)}(k)| \leq \Delta_C^{(l)} \wedge \mathbf{x}(k) \in \mathbf{X} \right\} \quad (16)$$

which is a strongly reduced subdomain of the original state domain \mathbf{X} in (6). In other words, the discretized SDP problem is solved in each iteration l within a corridor with width $\Delta_C^{(l)}$ around the received state trajectory $\mathbf{x}^{(l-1)}(k)$ of the previous iteration, to produce a solution trajectory $\mathbf{x}^{(l)}(k)$ for use in the next iteration. The corridor width $\Delta_C^{(l)}$, as well as the discretisation intervals $\Delta a^{(l)}, \Delta \mathbf{x}^{(l)}$, can vary in each iteration, typically at a decreasing rate. The iterations stops when a termination criterion is satisfied.

In the proposed GLOSA application (Typaldos et al., 2021), the starting trajectory for the first iteration of DDDP is the optimal solution of the deterministic GLOSA problem, assuming the "pessimistic" case where the traffic light will switch from red to green at the latest possible time, that is, at $k_1 = k_{max}$, so as to cover the whole time range and be not too far from the stochastic optimal trajectory; see (Typaldos et al., 2020a). The discretisation intervals $\Delta a^{(l)}, \Delta \mathbf{x}^{(l)}$ for the first iteration are also given. The intervals are reduced by half, each time there is no solution improvement at two subsequent iterations. The corridor width is taken as

$$\Delta_C^{(l)} = C \cdot \Delta a^{(l)} \quad (17)$$

where $C = [C_x, C_v]$ are constant given values specifying the state corridor width. Thus, the corridor's initial width is defined through the chosen values for C and $\Delta a^{(l)}$. In the following iterations, whenever the discrete interval is reduced, there is an analogous reduction of the corridor width. Consequently, we have a constant number of feasible discrete points in all iterations, which facilitates the algorithm's fine-tuning, so as to improve its computational efficiency. Note that, C_x and C_v values may differ, as the magnitude of the two respective state variables differs.

The admissible control region U (see (7)) is kept the same at each iteration, although many related state transitions cannot be considered in view of the reduced state variables domain. The algorithm terminates whenever there is no

improvement of the produced solution trajectory even after a reduction of the discretisation interval; or when the $\Delta a^{(l)}$ value becomes less than 0.125, which was found in (Typaldos et al., 2020a) to lead to sufficiently accurate results.

It should be noted that there is no general guarantee that the DDDP iterations will actually converge to the full-domain SDP solution. In particular, if the state sub-domains $\mathbf{X}_C^{(l)}$ considered in the iterations are too small, the obtained DDDP solution may actually differ from the overall SDP solution. On the other hand, if the state sub-domains are selected large, the required number of iterations may decrease, but the computation time required to find the solution at each iteration increases accordingly. In conclusion, some fine-tuning regarding the size of sub-domains is necessary to ensure convergence to the overall SDP solution with minimum overall (all iterations) computation time.

4. Differential Dynamic Programming (DDP)

DDP is an iterative algorithm, which requires a given initial feasible trajectory (of states and controls) to be used in the first iteration. At each iteration and time step, a quadratic-linear approximation of the recursive Bellman equation (14), around the initial trajectory of the iteration, is solved, leading to an improved trajectory. The improved trajectory is then used as initial trajectory in the next iteration, and the procedure continues until convergence to the solution of the original optimal control problem is achieved. The method was developed for deterministic optimal control problems; and it is not applicable to general stochastic problems, which do not feature solution trajectories (Bertsekas, 2012). However, as explained earlier, the specific stochastic GLOSA problem admits feasible solution trajectories for specific initial states, hence DDP may be applied.

4.1. General Constrained DDP

The standard DDP method (Jacobson and Mayne, 1970) does not account for constraints, so, for the needs of the problem presented in this work, extensions of DDP, proposed by Murray and Yakowitz (1979) for linear constraints and by Yakowitz (1986) for nonlinear constraints, is utilized.

Consider a discrete-time optimal control problem with states \mathbf{x} and controls \mathbf{a} with the recursive Bellman equation $V(\mathbf{x}(k), k) = \min\{\phi(\mathbf{x}, \mathbf{a}, k) + V(\mathbf{f}(\mathbf{x}, \mathbf{a}, k), k+1)\}$, for $k = K-1, \dots, 0$, where V is the optimal cost function, ϕ is the objective function of the optimal control problem, \mathbf{f} is the right-hand side of the state equation, and the minimisation must be carried out for all feasible controls $\mathbf{a}(k)$ that satisfy any inequality constraints present. The DDP procedure undertakes, at each time step of each iteration, a quadratic approximation of the term to be minimized in the recursive Bellman equation, i.e., of $\phi(\mathbf{x}, \mathbf{a}, k) + V(\mathbf{f}(\mathbf{x}, \mathbf{a}, k), k+1)$. The quadratic approximation is taken around nominal trajectories $(\bar{\mathbf{x}}(k), \bar{\mathbf{a}}(k))$, which are the initial trajectories of each iteration. Define $\delta\mathbf{x} = \mathbf{x} - \bar{\mathbf{x}}$ and $\delta\mathbf{a} = \mathbf{a} - \bar{\mathbf{a}}$. The aim is to find the optimal control law for $\delta\mathbf{a}$, which minimizes the quadratic approximation, subject to some inequality constraints. The procedure of the DDP algorithm for each iteration is described in brief as follows, see (Yakowitz, 1986; Murray and Yakowitz, 1979) for more details.

For $k = K-1, \dots, 0$, the quadratic approximation $Q(\mathbf{x}, \mathbf{a}, k)$ of the term $\phi(\mathbf{x}, \mathbf{a}, k) + V(\mathbf{f}(\mathbf{x}, \mathbf{a}, k), k+1)$ around a nominal point $(\bar{\mathbf{x}}(k), \bar{\mathbf{a}}(k))$, is required and can be expressed as follows (for simplicity, the indexes k are omitted)

$$Q(\mathbf{x}, \mathbf{a}) = \frac{1}{2}\delta\mathbf{x}^T \mathbf{A}\delta\mathbf{x} + \delta\mathbf{x}^T \mathbf{B}\delta\mathbf{a} + \frac{1}{2}\delta\mathbf{a}^T \mathbf{C}\delta\mathbf{a} + \mathbf{D}^T \delta\mathbf{a} + \mathbf{E}^T \delta\mathbf{x} \quad (18)$$

where the zero-order constant term is omitted, as it does not influence the minimisation; and $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{E}$ are matrices with appropriate dimensions and with elements that may depend on $(\bar{\mathbf{x}}(k), \bar{\mathbf{a}}(k))$.

Furthermore, we assume the general inequality constraints for the optimal control problem

$$\mathbf{c}(\mathbf{x}, \mathbf{a}) \leq \mathbf{0} \quad (19)$$

and we denote as $\tilde{\mathbf{c}}(\mathbf{x}, \mathbf{a})$ its linear approximation around the nominal trajectories $(\bar{\mathbf{x}}, \bar{\mathbf{a}})$.

For the solution of the resulting Quadratic Programming (QP) problem, Fletcher's active set algorithm is used (Fletcher, 1971). The critical issue here is that the minimization in the recursive Bellman equation should be executed in a parametrized way, i.e., deliver $\delta\mathbf{a}$ as a function of $\delta\mathbf{x}$, which is not compatible with a numerical solution procedure. Therefore, in the QP problem formulation, we replace \mathbf{x} with $\bar{\mathbf{x}}$ (i.e. we set $\delta\mathbf{x} = 0$) in (18) and $\tilde{\mathbf{c}}(\mathbf{x}, \mathbf{a})$, and solve the following QP problem to obtain a numerical value for $\delta\mathbf{a}$

$$\begin{aligned} & \min Q(\bar{\mathbf{x}}, \bar{\mathbf{a}} + \delta\mathbf{a}) \\ & \text{s.t} \\ & \tilde{\mathbf{c}}(\bar{\mathbf{x}}, \bar{\mathbf{a}} + \delta\mathbf{a}) \leq \mathbf{0} \end{aligned} \quad (20)$$

After solving the QP problem, the index set of the active constraints, denoted as S , leads to the linear system of equations

$$\tilde{c}_i(\bar{\mathbf{x}}, \bar{\mathbf{a}} + \delta \mathbf{a}) = 0, \quad i \in S \quad (21)$$

and expressing (21) in matrix form, we have

$$\mathbf{U} \delta \mathbf{a} - \mathbf{W} = \mathbf{0} \quad (22)$$

where \mathbf{U} has rank equal to the number of indexes in S . This solution of the QP problem satisfies the necessary conditions of optimality for the Lagrangian (Fletcher, 1971)

$$L(\mathbf{a}, \lambda) = \frac{1}{2} \delta \mathbf{a}^T \mathbf{C} \delta \mathbf{a} + \mathbf{D}^T \delta \mathbf{a} + \lambda^T (\mathbf{U}^T \delta \mathbf{a} - \mathbf{W}) \quad (23)$$

where λ is the Lagrange multiplier vector. These necessary conditions may be obtained by differentiating L with respect to $\delta \mathbf{a}$ and λ , which yields the following linear system of equations

$$\begin{bmatrix} \mathbf{C} & \mathbf{U} \\ \mathbf{U}^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \delta \mathbf{a} \\ \lambda \end{bmatrix} = \begin{bmatrix} -\mathbf{D} \\ \mathbf{W} \end{bmatrix} \quad (24)$$

In order to re-introduce $\delta \mathbf{x}$, i.e. the variability of state variables, we consider that the constraints that are active while we assumed $\mathbf{x} = \bar{\mathbf{x}}$, remain active for other \mathbf{x} , i.e.

$$\tilde{c}_i(\bar{\mathbf{x}} + \delta \mathbf{x}, \bar{\mathbf{a}} + \delta \mathbf{a}) = 0, \quad i \in S \quad (25)$$

which in matrix form reads

$$\mathbf{U} \delta \mathbf{a} - \mathbf{W} - \mathbf{X} \delta \mathbf{x} = \mathbf{0} \quad (26)$$

Then, we may derive a control law, i.e. $\delta \mathbf{a}$ as a function of $\delta \mathbf{x}$ by solving analytically the accordingly extended version of (24), i.e.

$$\begin{bmatrix} \mathbf{C} & \mathbf{U} \\ \mathbf{U}^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \delta \mathbf{a} \\ \lambda \end{bmatrix} = \begin{bmatrix} -\mathbf{D} - \mathbf{B} \delta \mathbf{x} \\ \mathbf{W} + \mathbf{X} \delta \mathbf{x} \end{bmatrix} \quad (27)$$

which yields the following linear control law with readily computed vector α and matrix β

$$\delta \mathbf{a}(\mathbf{x}) = \alpha + \beta \delta \mathbf{x} \quad (28)$$

This minimisation outcome is substituted in the quadratic approximation Q to obtain, from the Bellman equation, the approximate optimal cost function

$$V(\mathbf{x}(k), k) = Q(\mathbf{x}(k), \alpha(k) + \beta(k) \delta \mathbf{x}(k), k) \quad (29)$$

Which must be stored for use at the next time step $k + 1$ within the same iteration. Note that this approximate cost function is quadratic by construction, which facilitates the quadratic approximation in the Bellman equation at the next time step $k + 1$.

Given $\delta \mathbf{a}(\delta \mathbf{x}, k)$, from (28), the nominal trajectories for the next iteration are calculated as follows for $k = 0, \dots, K-1$

$$\mathbf{a}(k) = \bar{\mathbf{a}}(k) + \epsilon [\alpha(k) + \beta(k) \delta \mathbf{x}(k)] \quad (30)$$

$$\mathbf{x}(k+1) = f(\mathbf{x}(k), \mathbf{a}(k), k) \quad (31)$$

$$\mathbf{x}(0) = \mathbf{x}_0 \quad (32)$$

where $0 < \epsilon \leq 1$ is used to ensure the convergence of the algorithm. Specifically, ϵ values smaller than 1 may facilitate convergence if needed; but may decrease the convergence rate if selected too small. The algorithm terminates whenever the following convergence test is satisfied

$$\left(\sum_{k=0}^{K-1} [(\mathbf{a}(k) - \bar{\mathbf{a}}(k))^2] \right)^{1/2} < \epsilon_1 \quad (33)$$

where ϵ_1 is a small positive number.

In contrast to the SDP and DDDP algorithms, DDP does not apply any discrete DP procedures, hence no discretization of the state and control domains is necessary. As a result, the required computation time is independent of any discretization intervals, and the obtained solution trajectories are value-continuous, hence potentially better than those obtained by the other algorithms, even when using fine discretization.

4.2. Constrained DDP for the GLOSA Problem

To apply the DDP algorithm to the GLOSA stochastic optimal control problem described in Section 2.2, the quadratic approximation of the right-hand side of (14) around a nominal trajectory $(\bar{\mathbf{x}}(k), \bar{a}(k))$ is first required. Note that only the J_{DG}^* term needs to be approximated, as the rest terms are already quadratic. For the approximation, the LOESS (locally estimated scatterplot smoothing) method was utilized (Cleveland et al., 2017). Specifically, a number of (x, v) points are specified around the nominal trajectory at each iteration, and the corresponding values of $J_{DG}^*(\mathbf{x}(k), kT)$ are computed. It was found that some 4 such points per time step are sufficient for a good fit. Then, LOESS fits the following smooth function to the data set

$$p(x, v) = p_1 \delta x^2 + p_2 \delta v^2 + p_3 \delta x \delta v + p_4 \delta x + p_5 \delta v + p_6 \quad (34)$$

where p_1, \dots, p_6 are the coefficients of fitting. The resulting fitted function (34), along with the other quadratic terms of (14), can be easily expressed in the form of (18). It was verified that the resulting quadratic function (18) is always convex, something that was expected from the physical background of the problem.

The nominal trajectory of the first iteration is taken to be the optimal solution of the deterministic GLOSA problem, assuming the "pessimistic" case where the traffic light switches from red to green at the latest possible time, i.e., $k_1 = k_{max}$.

For the state constraints (6), we may replace the state equations (4), (5) therein, hence, for the admissible region of the SDP (6) and (7), we have the equivalent constraints

$$\begin{aligned} \frac{2(x_{min} - x(k) - v(k)T)}{T^2} &\leq a(k) \leq \frac{2(x_{max} - x(k) - v(k)T)}{T^2} \\ \frac{v_{min} - v(k)}{T} &\leq a(k) \leq \frac{v_{max} - v(k)}{T} \\ a_{min} &\leq a(k) \leq a_{max} \end{aligned} \quad (35)$$

which can be written in terms of δx and δa as follows

$$\begin{bmatrix} 1 \\ -1 \\ 1 \\ -1 \\ 1 \\ -1 \end{bmatrix} \delta a - \begin{bmatrix} \frac{2(x_{max} - \bar{x}(k+1))}{T^2} \\ -\frac{2(x_{min} - \bar{x}(k+1))}{T^2} \\ \frac{(v_{max} - \bar{v}(k+1))}{T} \\ -\frac{(v_{min} - \bar{v}(k+1))}{T} \\ \frac{(a_{max} - \bar{a}(k))}{T} \\ -\frac{(a_{min} - \bar{a}(k))}{T} \end{bmatrix} - \delta \mathbf{x} \leq \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (36)$$

Having the quadratic approximation of the recursive Bellman equation and the linear inequality constraints (35), the resulting QP is as described in the previous subsection. In fact, the resulting QP problem for the stochastic GLOSA application is simpler than the general case of Section 4.1. After setting $\delta \mathbf{x} = 0$ in (36), we have only one decision variable, δa , with constant upper and lower bounds. Thus, for the solution of the QP problem, it suffices to first compute analytically the value of δa that minimises the quadratic function. If this value does not violate any of the upper or lower bounds, then the QP solution is given by that value; else, the QP solution equals the activated bound. All further procedures are as explained in Section 4.1.

5. Demonstration Results and Comparison

In this section, the performance of the proposed approaches is tested, demonstrated and compared. Specifically, the DDDP and DDP algorithms are applied to realistic GLOSA scenarios and evaluated regarding the quality of their

Table 1

Investigated scenarios.

	x_0 (m)	v_0 (m/s)	x_e (m)	v_e (m/s)	x_1 (m)
Scenario 1	0	5	220	11	150
Scenario 2	0	11	220	11	150
Scenario 3	50	11	220	11	150

solution compared to the standard SDP, as well as the computational time needed for each algorithm to converge to the optimal solution.

5.1. Scenarios Setup

For the evaluation, several GLOSA-relevant scenarios have been considered in various situations, such as cases where the vehicle needs to accelerate or decelerate before or after the traffic signal; or cases where the vehicle is forced to fully stop and wait until the traffic light switches from red to green. Note that in all scenarios the aim is for a single vehicle to cross an intersection without considering interactions in case of slower preceding vehicles. However, the procedure in case of impending traffic is quite straightforward and is discussed in Section 6. For this paper, three scenarios have been chosen for demonstration with full illustration of state and control trajectories, according to Table 1. For consistency, the chosen scenarios are the same as in previous related works (Typaldos et al., 2020a, 2021). All scenarios concern mixed acceleration-deceleration cases. However, in Scenario 1, the vehicle starts with low initial speed and needs to mostly accelerate in order to optimally pass the traffic signal. On the other hand, in Scenarios 2 and 3, the vehicle starts at higher speed, hence it needs to mostly decelerate to avoid a full stop at the traffic signal. Compared to Scenarios 1 and 2, in Scenario 3 the vehicle starts closer to the traffic light, which leads to a full stop for few time steps before the traffic signal switches from red to green.

For all evaluated scenarios, the time step T is 1 s and the weight parameter w in (3) is equal to 0.1. The bounds for the states and control are set to $[x_{min}, x_{max}] = [0, 150]$ m, $[v_{min}, v_{max}] = [0, 16]$ m/s and $[a_{min}, a_{max}] = [-3, 3]$ m/s², respectively. For the traffic signal parameters, we assume that the certain red phase is $k \in [0, k_{min} - 1 = 9]$ s, while the window of possible switching times, from red to green, is $[k_{min}, k_{max}] = [10, 30]$ s with uniform a-priori probability distribution. Regarding DDDP, for the initial discretization, $\Delta a = \Delta v = 0.5$ is used, which leads to $\Delta x = 0.25$ m according to (15); and the initial corridor width in (17) is set $C = [20, 4]$, i.e., we have $C_x = 5 \cdot C_v$. The choice of the initial discretization and corridor width value will be justified later in this section. For the DDP algorithm, parameters ϵ in (30) and ϵ_1 in (33) are set equal to 1 and 10^{-4} , respectively.

5.2. DDDP algorithm results

Fig. 2 (Scenario 1) and Figs. 3 and 4 (for Scenarios 2 and 3, respectively) display the optimal state (speed) and control (acceleration) trajectories over each iteration of the DDDP algorithm. Note that, position trajectories are not included, as the differences in those trajectories, at each iteration, are small and barely visible. In each iteration, we consider a corridor $\Delta_C^{(l)} = [-C \cdot \Delta a^{(l)}, C \cdot \Delta a^{(l)}]$ around the respective received state trajectories, which cannot of course extend out of the full state bounds.

In all Figs 2, 3 and 4, the dashed blue lines represent, for each iteration, the received trajectory, the solid orange lines represent the optimal trajectories derived, and the red dashed lines reflect the corresponding corridor bounds. For Scenario 1 (Fig. 2), it can be observed that, starting with the initial chosen discretization, the first DDDP iteration improves the initial trajectory, leading to a better solution, which is optimal within the considered sub-domain. In the second iteration, no further improvement can be achieved within the new sub-domain bounds, which means that, with the current discretization values, the best achievable solution has been reached. By reducing the discretization at the 3rd iteration, a reduction in the corridor width is observed, but the number of discrete points remains the same. This reduction enables further improvement of the solution, which remains the same in iteration 4; hence the discretisation is further reduced, leading to further improvement in iteration 5; and the procedure continues with iteration 6, where no further improvement is reached and the algorithm terminates, as the discretisation has reached the minimum required value. Similar behaviour is observed for Scenarios 2 and 3, where 10 iterations are required for convergence, and the respective Figs. 3 and 4 display a subset thereof. Table 2 contains the values of the objective function for each iteration of the DDDP algorithm for all scenarios, assuming $\Delta a = 0.5$ and $C = [20, 4]$, where the choice of these values will be explained in the following.

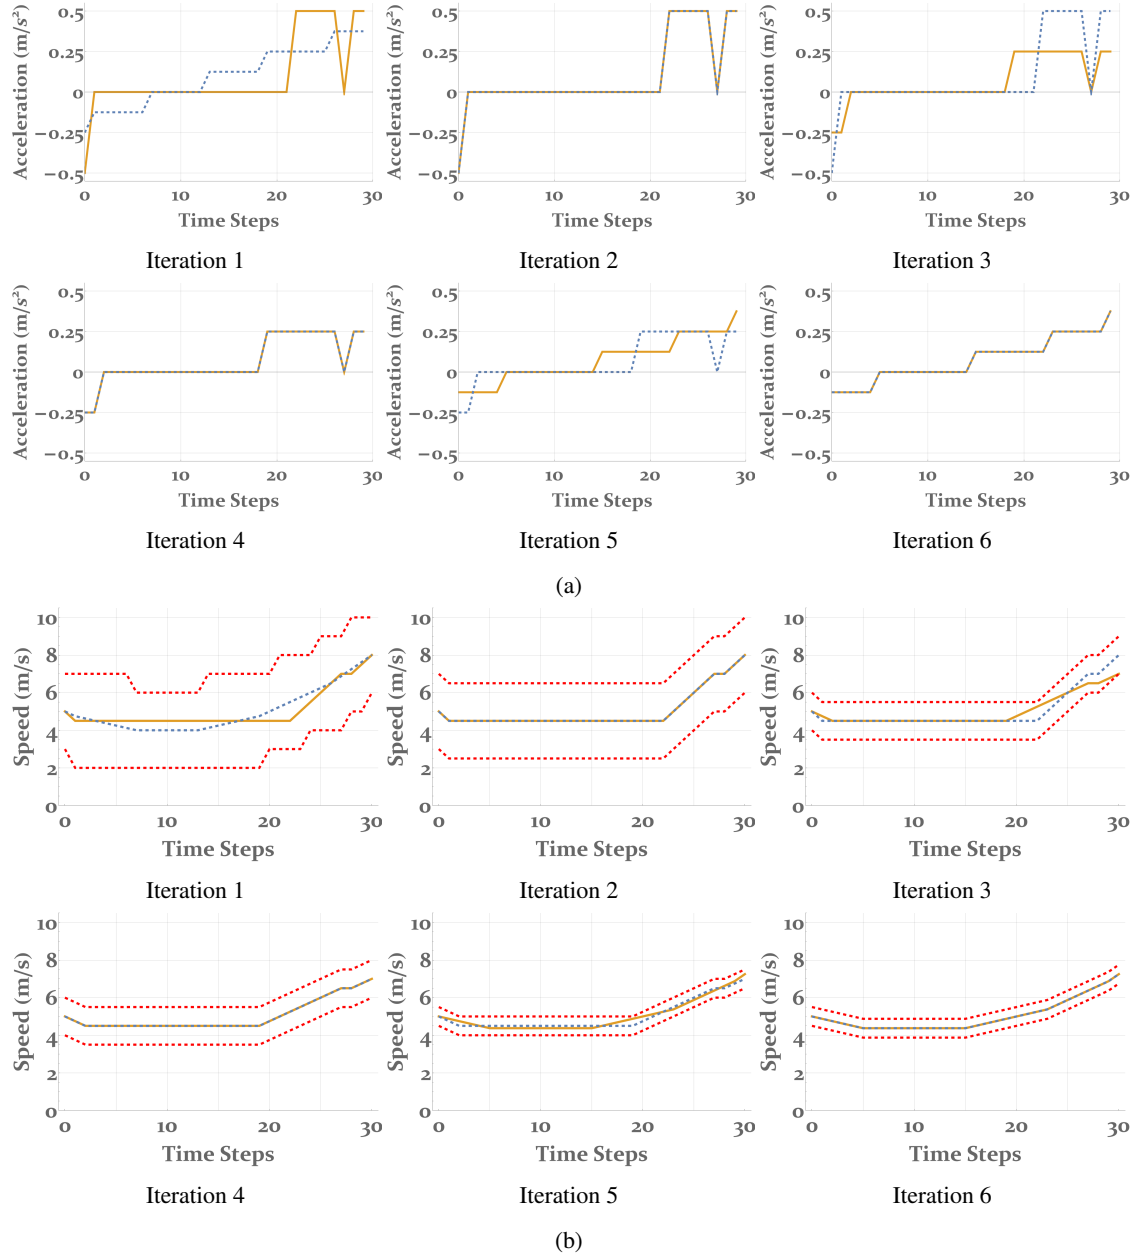


Figure 2: Received (blue dashed line) and optimal (orange line) (a) acceleration and (b) speed trajectories of DDDP algorithm in each iteration for Scenario 1.

The DDDP has been evaluated for different corridor C_v and initial discretization Δa values, for the three investigated scenarios (see Table 3). It can be seen that, for all presented initial discretization values, the obtained optimal solution at convergence reaches the same globally optimal solution, if the corridor width is sufficiently large, specifically for $C_v \geq 4$ (or lower, in some cases). This is because, for very small values of C_v , DDDP does not have sufficient flexibility to improve the solution at each iteration due to the extremely limited state space, something that may lead to a local optimum of the overall problem at convergence. At the same time, a reduction of the number of DDDP iterations is obtained for larger corridors. This behaviour is expected, as the bigger corridors, i.e., bigger admissible sub-regions, lead to potentially better solutions at each iteration, and hence to fewer iterations to converge to the optimal solution. On the other hand, despite the decrease on the number of iterations, the overall computation time may increase due

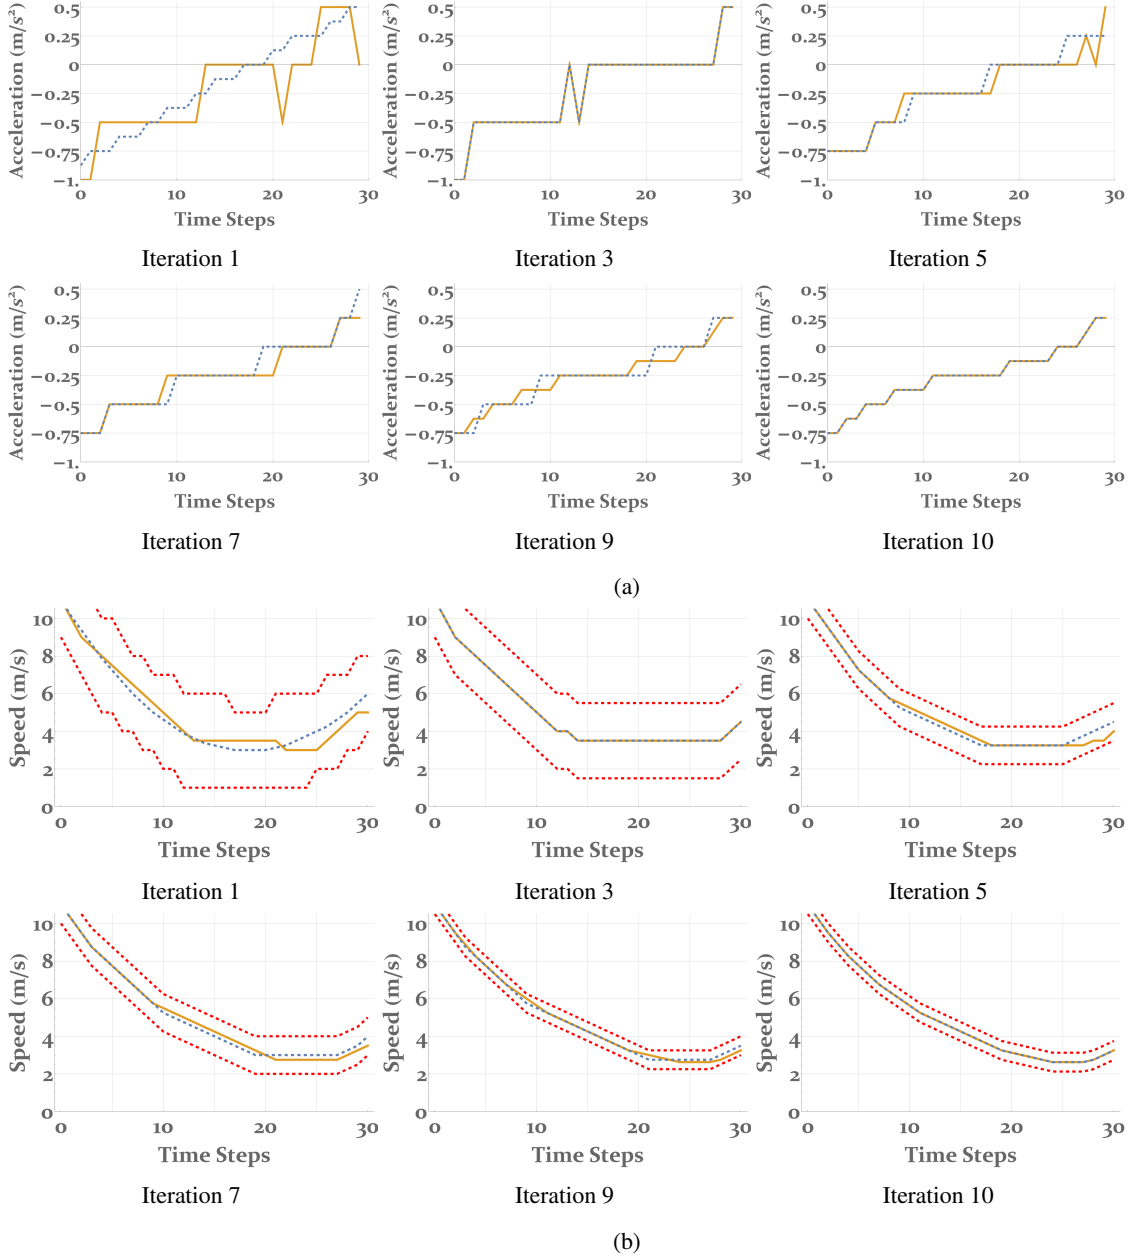


Figure 3: Received (blue dashed line) and optimal (orange line) (a) acceleration and (b) speed trajectories of DDDP algorithm in each iteration for Scenario 2.

to the higher computation time required at each iteration, which in turn, is due to more feasible discrete state points included in each iteration's sub-region. Based on these observations, the selection of the values $C = [20, 4]$ and $\Delta a = 0.5$ seems to be a reasonable choice.

The accuracy of the DDDP algorithm compared to the full one-shot SDP solution may be assessed by contrasting the optimal cost and the optimal states and control trajectories obtained in the last respective DDDP iteration of the three scenarios, to the corresponding optimal trajectories derived from the full-range SDP with a discretisation of $\Delta a = 0.125$. As a matter of fact, DDDP converges, in all cases, to the exact same optimal solution as the full-range SDP. The obtained optimal costs of both approaches, for the three scenarios, are 1.175, 3.906 and 6.3582, respectively, while the computation time differences are remarkable, as the one-shot SDP needs approximately 613 s to obtain the

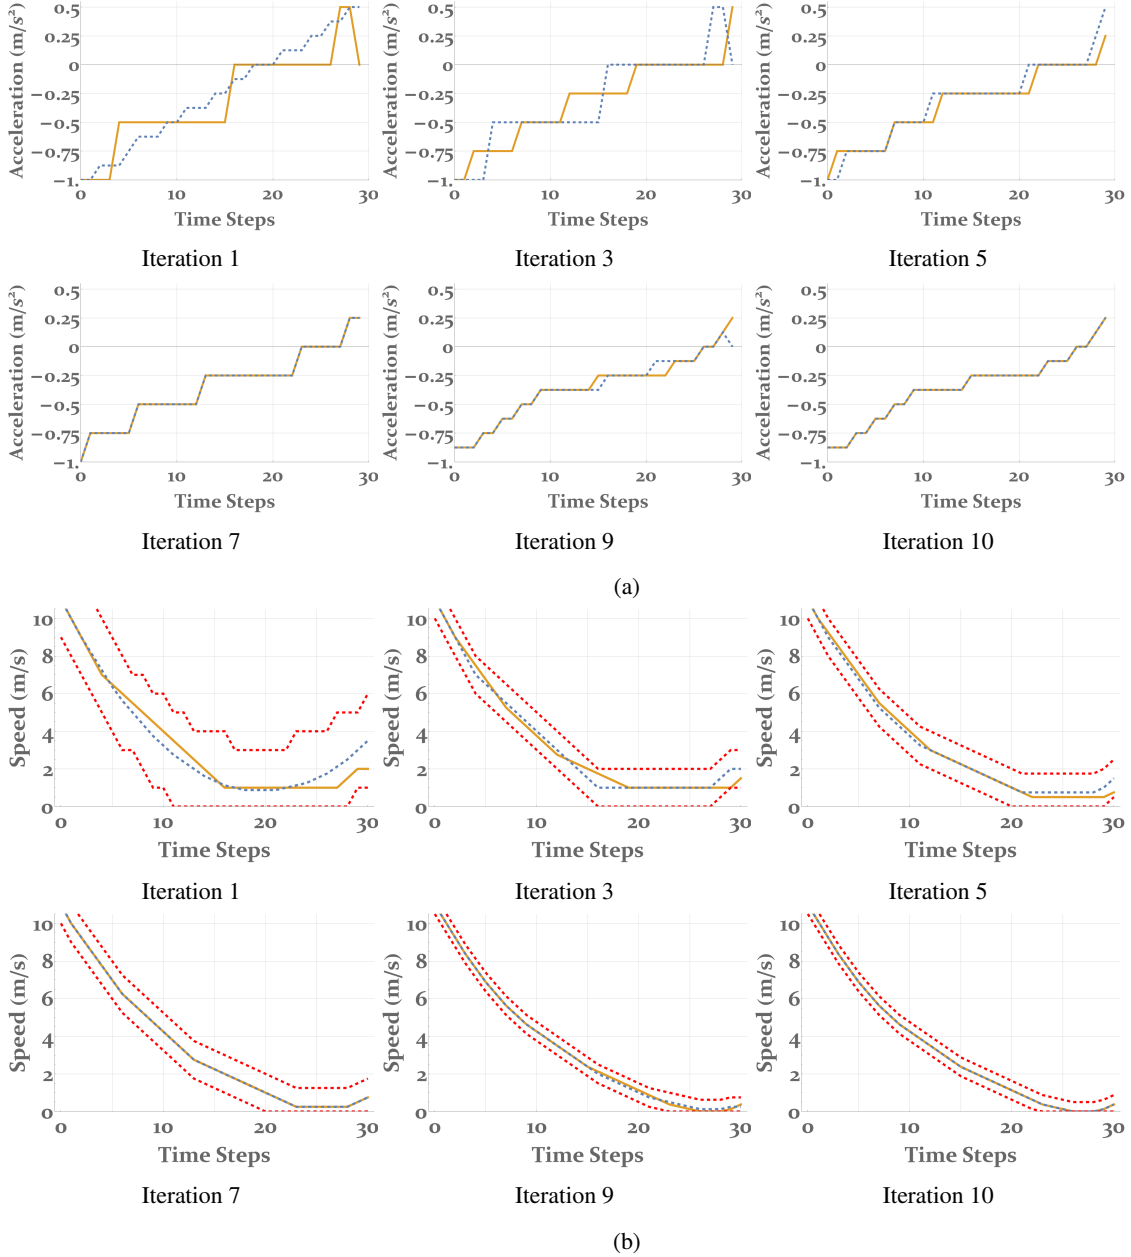


Figure 4: Received (blue dashed line) and optimal (orange line) (a) acceleration and (b) speed trajectories of DDDP algorithm in each iteration for Scenario 3.

solution (same for all scenarios), while DDDP needs only 0.69, 1.24 and 1.12 s for the respective three scenarios. More importantly, this big reduction in computation time enables the DDDP algorithm to be executable in real time, even in an MPC mode, on the vehicle side, similarly to the deterministic GLOSA.

5.3. DDP algorithm results

Fig. 5 (Scenario 1) and Figs. 6 and 7 (for Scenarios 2 and 3, respectively) illustrate the evolution of the state (speed) and control (acceleration) trajectories produced by the DDP algorithm at each iteration. In all figures, blue dashed lines represent the initial (received) trajectories, and solid orange lines the derived trajectories of each iteration. From the figures, it can be seen that the DDP algorithm, starting from the pessimistic deterministic GLOSA solution,

Table 2

Optimal cost evolution and discretization change in each iteration of DDDP algorithm.

Scenario 1				Scenario 2			Scenario 3		
Iter.	$\Delta a = \Delta v$	Δx	Cost	$\Delta a = \Delta v$	Δx	Cost	$\Delta a = \Delta v$	Δx	Cost
1	0.5	0.25	1.357	0.5	0.25	4.311	0.5	0.25	6.6218
2	0.5	0.25	1.357	0.5	0.25	4.231	0.5	0.25	6.6218
3	0.25	0.125	1.223	0.5	0.25	4.231	0.25	0.125	6.4532
4	0.25	0.125	1.223	0.25	0.125	4.012	0.25	0.125	6.4346
5	0.125	0.0625	1.175	0.25	0.125	3.988	0.25	0.125	6.4208
6	0.125	0.0625	1.175	0.25	0.125	3.958	0.25	0.125	6.4091
7				0.25	0.125	3.940	0.25	0.125	6.4091
8				0.25	0.125	3.940	0.125	0.0625	6.3583
9				0.125	0.0625	3.906	0.125	0.0625	6.3582
10				0.125	0.0625	3.906	0.125	0.0625	6.3582

Table 3Performance of DDDP algorithm in terms of CPU-time and optimal cost for different initial values of C and Δa .

Scenario 1												
$\Delta a = 1.0$				$\Delta a = 0.5$			$\Delta a = 0.25$			$\Delta a = 0.125$		
C_v	Iter.	CPU Times (s)	Cost	Iter.	CPU Times (s)	Cost	Iter.	CPU Times (s)	Cost	Iter.	CPU Times (s)	Cost
2	19	0.648	1.199	20	0.612	1.193	10	0.326	1.193	5	0.280	1.193
3	19	1.296	1.177	8	0.552	1.175	6	0.469	1.179	5	0.443	1.179
4	15	1.655	1.175	6	0.692	1.175	7	0.867	1.175	5	0.762	1.175
5	12	1.839	1.175	6	1.052	1.175	6	1.111	1.175	4	0.885	1.175
6	10	2.298	1.175	6	1.451	1.175	5	1.291	1.175	3	0.967	1.175
Scenario 2												
$\Delta a = 1.0$				$\Delta a = 0.5$			$\Delta a = 0.25$			$\Delta a = 0.125$		
C_v	Iter.	CPU Times (s)	Cost	Iter.	CPU Times (s)	Cost	Iter.	CPU Times (s)	Cost	Iter.	CPU Times (s)	Cost
2	29	1.212	3.909	18	0.790	3.909	16	0.729	3.909	21	1.014	3.909
3	22	1.656	3.906	14	1.093	3.906	12	0.966	3.906	18	1.568	3.906
4	13	1.578	3.906	10	1.246	3.906	9	1.149	3.906	13	1.740	3.906
5	11	1.920	3.906	7	1.273	3.906	8	1.514	3.906	11	2.178	3.906
6	14	3.324	3.906	10	2.471	3.906	7	1.823	3.906	9	2.457	3.906
Scenario 3												
$\Delta a = 1.0$				$\Delta a = 0.5$			$\Delta a = 0.25$			$\Delta a = 0.125$		
C_v	Iter.	CPU Times (s)	Cost	Iter.	CPU Times (s)	Cost	Iter.	CPU Times (s)	Cost	Iter.	CPU Times (s)	Cost
2	28	1.152	6.3583	20	0.868	6.3583	18	0.815	6.3583	21	0.998	6.3644
3	21	1.543	6.3582	16	1.192	6.3582	14	1.098	6.3582	17	1.419	6.3639
4	14	1.506	6.3582	10	1.128	6.3582	10	1.179	6.3582	14	1.834	6.3582
5	12	1.777	6.3582	8	1.253	6.3582	8	1.400	6.3582	11	2.101	6.3582
6	11	2.050	6.3582	7	1.368	6.3582	7	1.553	6.3582	9	2.328	6.3582

converges to the overall optimal solution. It can also be observed that, at each iteration, starting from the initial trajectory, the algorithm results in an improved one, corresponding to an improved solution. The DDP algorithm is seen to converge very fast to the optimal solution, with the improvement being more profound at the first iterations, as the change between the nominal and the produced trajectories are more noticeable. On the other hand, in the last iterations, the improvement is smaller, as the algorithm is already very close to the optimal solution. As mentioned

earlier, all intermediate DDP solutions at each iteration are continuous due to the nature of DDP, which is operating on a continuous-value basis for all involved variables.

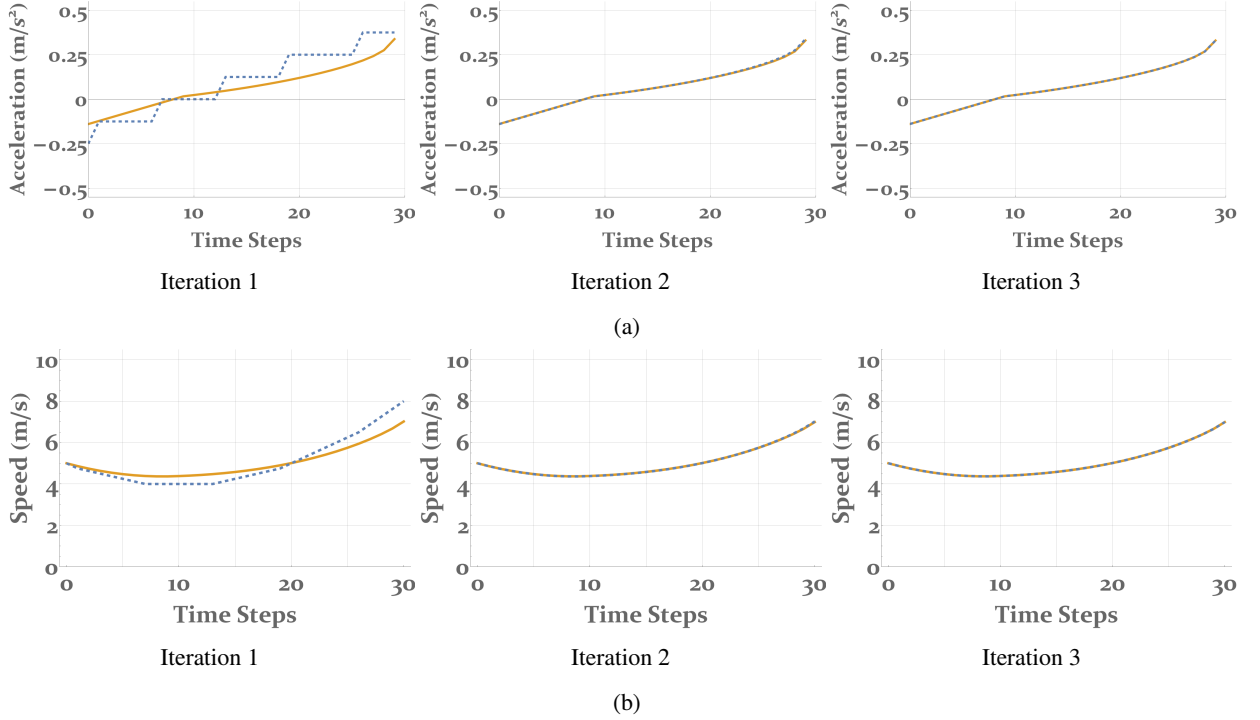


Figure 5: Received (blue dashed line) and computed (orange line) trajectories of (a) acceleration and (b) speed at each iteration of DDP algorithm for Scenario 1.

Fig. 8 demonstrates the accuracy of the DDP solutions compared to the one-shot discrete SDP (or DDDP, as it results in the same optimal trajectories as the standard SDP algorithm) solutions for the three scenarios. In this figure, the optimal state (speed) and control (acceleration) trajectories obtained from DDP algorithm at convergence are contrasted to those obtained from the full-range SDP algorithm with discretization $\Delta a = 0.125 \text{ m/s}^2$. The two solutions, in all scenarios, are very close to each other, with a slight advantage for the DDP solution, which is not bound to a fine, but still visible value-discretization of the state and control, as it is the case for the one-shot discrete SDP algorithm.

Fig. 8 also includes the analytical solution of the respective deterministic GLOSA problems (dashed lines) after the switching from red to green, which is assumed here to occur at the latest possible time, i.e., at k_{max} . As mentioned in Section 2.1, the deterministic GLOSA problem solution depends on the initial state of the vehicle and the switching time. Note that, for a given intersection, the final state is the same for any approach. Therefore, while the continuity between the optimal state trajectories of the numerical and the analytical solutions is not an issue, as the latter is just a continuation of the former, the change in acceleration trajectories is quite small as it is verified from Fig. 8.

Finally, we note (without providing figures) that the jerk trajectories of both stochastic and deterministic GLOSA solutions, for both DDP and the one-shot SDP, exhibit low values, including the point of switching between the two trajectories, thus ensuring passenger comfort.

5.4. Summarized results

Table 4 contains the summarized information of the number of iterations until convergence, CPU time and the optimal obtained cost of the one-shot SDP, the iterative DDDP and the iterative but non-discrete DDP algorithms for Scenarios 1, 2 and 3. Based on the results reported in the table, it can be verified that DDP's optimal solution is slightly better compared to the other two approaches due to considering non-discretised control trajectories. However, the most striking difference of DDP is the computational time needed to solve the problem, which is extremely small. Specifically, the one-shot SPD takes 631 s, but it needs to be solved just once to deliver the optimal closed-loop control

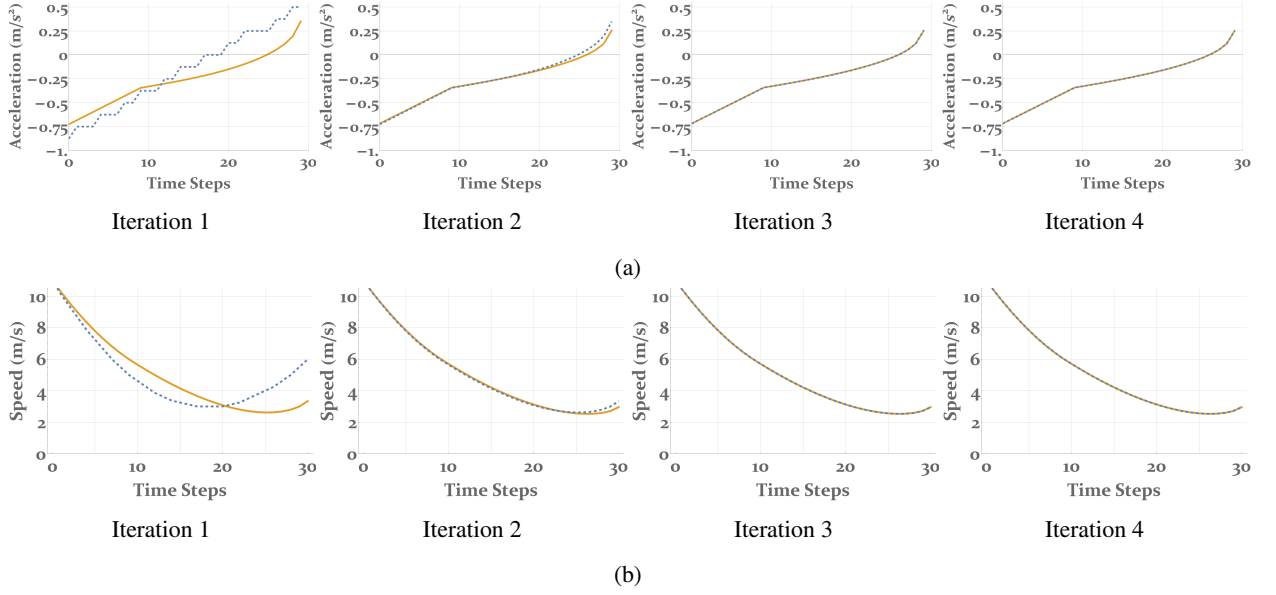


Figure 6: Received (blue dashed line) and computed (orange line) trajectories of (a) acceleration and (b) speed at each iteration of DDP algorithm for Scenario 2.

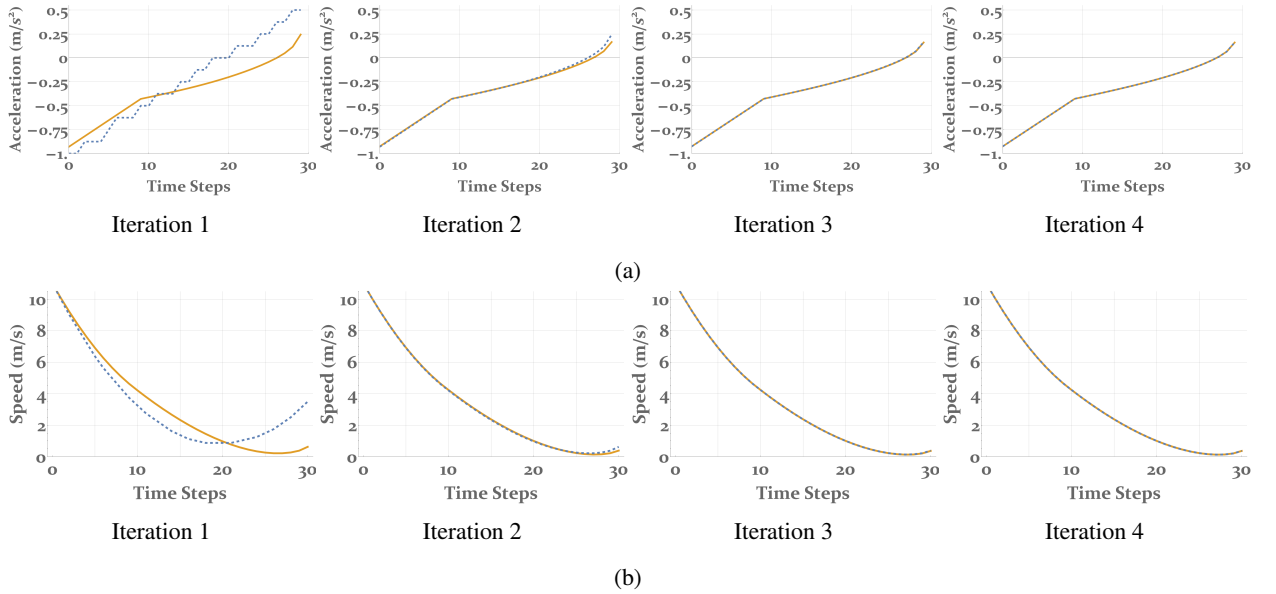


Figure 7: Received (blue dashed line) and computed (orange line) trajectories of (a) acceleration and (b) speed at each iteration of DDP algorithm for Scenario 3.

law (see also Section 2.3) which applies to any possible scenario. On the other hand, DDDP significantly reduces the required CPU time with the solutions being derived in around 1 s, while DDP achieves quasi-instantaneous solutions in all scenarios.

In order to explore the robustness of the DDDP and DDP algorithms and to further emphasize their superiority compared to the one-shot SDP, in terms of computation times, a more excessive analysis was conducted and is presented in Fig. 9. Specifically, Fig. 9 illustrates the required CPU times until convergence for both proposed algorithms for a plethora of different initial conditions, i.e., initial speeds and positions (distance from the traffic signal) of a vehicle. From both Fig. 9a and Fig. 9b it can be seen that, for all evaluated initial conditions, the CPU times of both

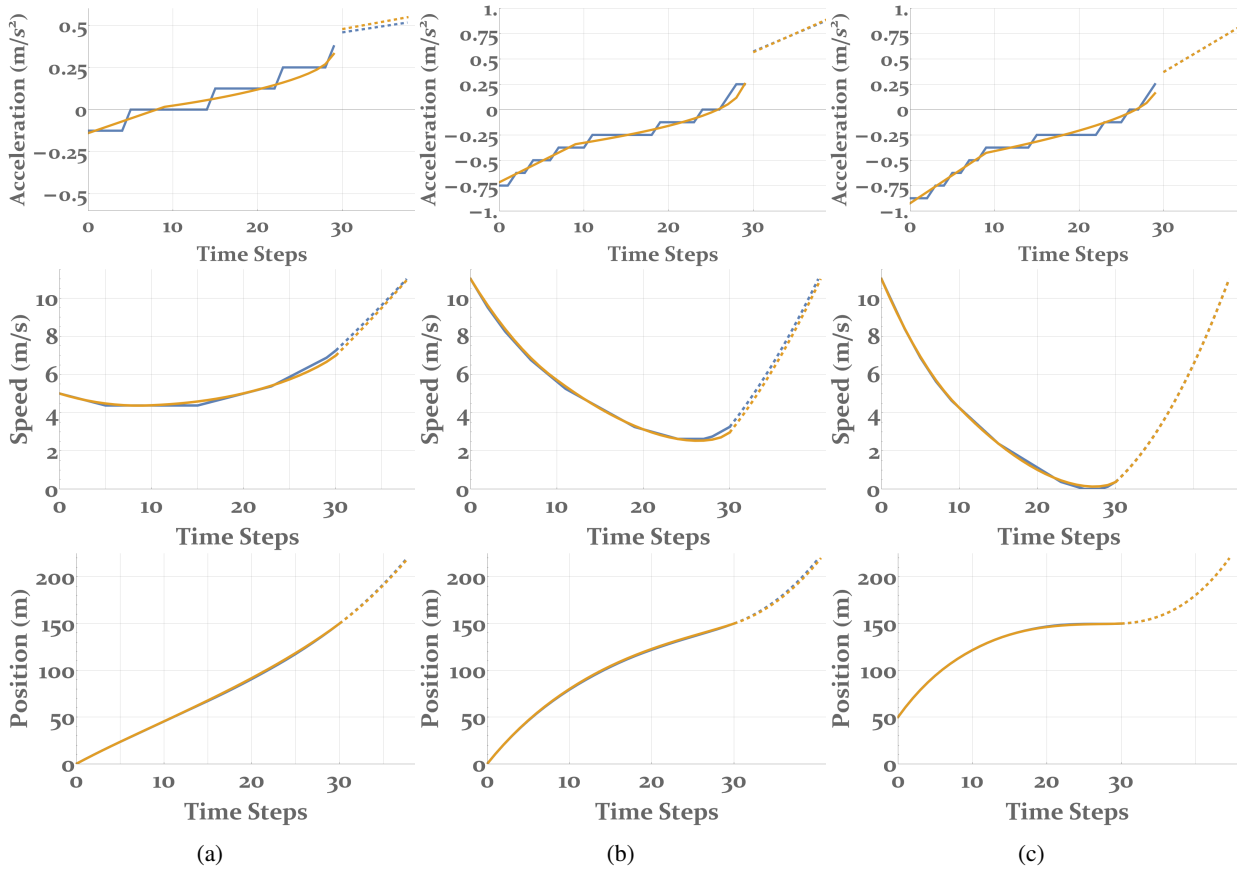


Figure 8: Optimal acceleration, speed and position trajectories of the last DDP iteration (orange solid lines) and the one-shot SDP (blue solid lines), along with the corresponding deterministic-GLOSA follow-up solutions (orange and blue dashed lines, respectively) for (a) Scenario 1, (b) Scenario 2 and (c) Scenario 3.

Table 4

Comparison of iterations, CPU times until convergence and obtained optimal cost of SDP, DDDP and DDP

Scenario No.	One-shot SDP			DDDP			DDP		
	1	2	3	1	2	3	1	2	3
Iterations	-			6	10	10	3	4	4
CPU Times (s)	631 (for all possible scenarios)			0.692	1.246	1.128	0.0005	0.0007	0.0007
Cost	1.175	3.906	6.358	1.175	3.906	6.358	1.162	3.892	6.353

DDDP and DDP are similarly low as for the above three scenarios, with average times of 1.4 s and 0.001 s, respectively, the corresponding maximum and minimum times being 6.3 s and 0.2 s for DDDP and 0.006 s and 0.0005 s for DDP, respectively. These findings lead to the conclusion that, even though DDDP reduces the CPU time significantly compared to the one-shot SDP, the DDP algorithm achieves almost instantaneous solutions, which enables the DDP approach to be used, within an MPC framework, in the vehicle's on-board computer.

6. Discussion of Features, Limitations and Usage

For the standard SDP algorithm, as it is known, the complexity increases exponentially with the number of states and controls, whereby the basis of the exponential term is the number of considered discrete values, which increases

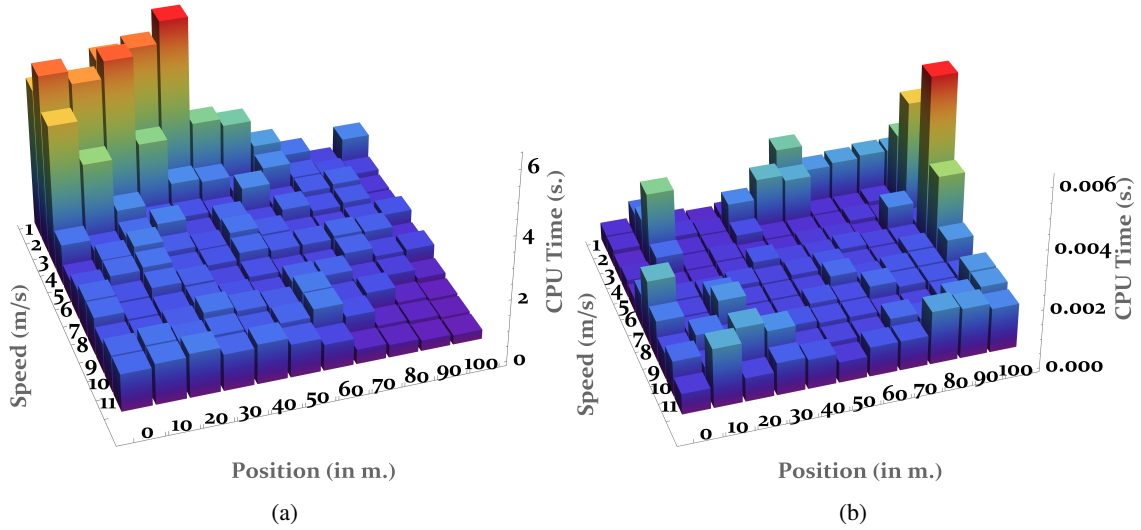


Figure 9: Computation times until convergence of: (a) DDDP and (b) DDP algorithms for different initial conditions.

linearly when the admissible region is increasing (Papageorgiou et al., 2015). On the other hand, DDDP algorithm's performance is mainly affected by the width of the corridor around the initial guess, the initial discretization level (see Table 3) and how good the initial guess may be in relation to the optimal solution. Regarding DDP, the algorithm solves, at each iteration and time step, a quadratic-linear approximation of the recursive Bellman equation around an initial trajectory. Thus, the algorithm's computational performance depends mainly on the quality of the initial trajectory of the first iteration and the quality of the approximation of the recursive Bellman equation.

The GLOSA problem, and hence all reported solution approaches, assumes that there is no interaction with front vehicles, which is also the case for most available GLOSA apps. Thus, in the case of slower preceding vehicles, the GLOSA problem (deterministic and stochastic alike) assumptions are not met. A possible procedure that may be applied in such cases is as follows (see also Fig. 10). The vehicle obtains the optimal GLOSA trajectories; however, the acceleration that the vehicle actually applies is the minimum between the GLOSA-optimal and the car-following (or ACC) accelerations. This procedure can be applied repeatedly, in an MPC mode, which is feasible thanks to the low CPU times.

Moreover, this work considers one type of vehicles, i.e., conventional gasoline-powered vehicles, driving on a straight roadway with zero grade. However, the basic SDP formulation and the solution algorithms presented in this paper allow for problem modifications for consideration of important special cases, such as roadway grades (positive or negative), different vehicle models (light-duty or heavy-duty vehicles, gasoline or diesel engines, etc.), battery electric vehicles, hybrid electric vehicles, hydrogen vehicles etc. This is also facilitated by the fact that the present work enables on-board solutions, which implies that the problem solved in each vehicle can be adapted to the characteristics of that vehicle and of the forthcoming intersection (e.g., in case of grade). However, the use of complex fuel-consumption or vehicle models may render the analytical solution of the deterministic GLOSA problem more difficult, in which case appropriate approximations may be required, e.g., quadratic approximations of a complex nonlinear objective function, as proposed in (Typaldos et al., 2020b).

The present work concentrated on the methodological and challenging algorithmic aspects of the stochastic GLOSA problem, while other practical aspects, such as V2I communication or experimental results with real data, as well as the impact of slower downstream traffic (as illustrated in Fig. 10) are left for future investigations. It should be emphasised that our method allows for any probability distribution of signal switching times to be considered; and to be changed over time, since using different probabilistic distributions does not alter neither the stochastic problem, nor the algorithms, nor the execution times.

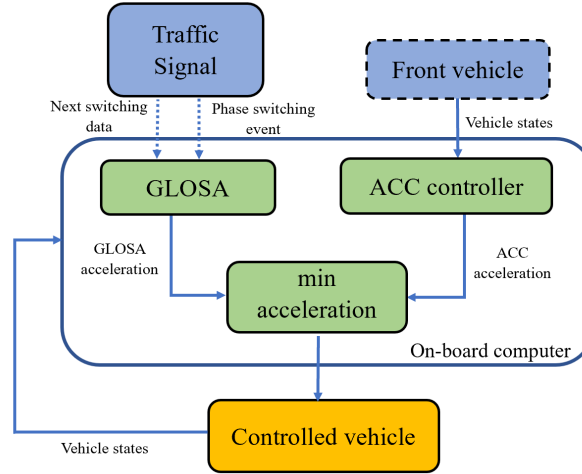


Figure 10: Procedure in case of impending traffic. Dashed line arrows represent interactions needed only once, while vehicle states correspond to vehicle's position and speed.

7. Conclusions

In recent works (Typaldos et al., 2020a, 2021), a stochastic GLOSA methodology was developed, by optimizing, using SDP and DDDP techniques, respectively, the vehicle kinematic trajectories subject to the intermediate stochastic traffic signal switching constraint and with a fixed final state. As an extension to these works, an iterative Differential Dynamic Programming (DDP) algorithm was employed in this paper. Within each iteration, DDP solves, for each time-step separately, a quadratic-linear approximation of the recursive Bellman equation. Demonstration and comparison results indicate that the DDP algorithm strongly outperforms both the full-range SDP and the DDDP algorithms in terms of computation time; and even delivers a slightly better optimal solutions due to lack of discretisation. This enables the DDP algorithm to be readily executable within a model predictive control framework in real time on-board approaching vehicles.

Possible ways of using GLOSA procedures in cases of impending traffic were outlined, but the scenario-based investigation of such procedures or of many-intersection crossing, which are highly interesting, are left for future work. Further current and future work is focused on:

- Generalization of the current GLOSA problem by considering uncertain switching times for both green and red phases.
- Testing of the approach performance in presence of multiple (equipped and not equipped) vehicles, while traversing multiple subsequent intersections.

Acknowledgments

The research leading to these results has received funding from the European Research Council under the European Union's Horizon 2020 Research and Innovation programme/ERC Grant Agreement n. [833915], project TrafficFluid.

References

- Bertsekas, D., 2012. Dynamic programming and optimal control: Volume I. volume 1. Athena scientific.
- Blokpoel, R., Lu, M., 2017. Signal plan stabilization to enable eco-driving, in: 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC), IEEE. pp. 133–138.
- Bodenheimer, R., Brauer, A., Eckhoff, D., German, R., 2014. Enabling glosa for adaptive traffic lights, in: 2014 IEEE vehicular networking conference (vnc), IEEE. pp. 167–174.
- Budhiraja, R., Carpentier, J., Mastalli, C., Mansard, N., 2018. Differential dynamic programming for multi-phase rigid contact dynamics, in: 2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids), IEEE. pp. 1–9.
- Cleveland, W.S., Grosse, E., Shyu, W.M., 2017. Local regression models, in: Statistical models in S. Routledge, pp. 309–376.

- Eteifa, S., Rakha, H.A., Eldardiry, H.M., Center, E., et al., 2021. Estimating Switching Times of Actuated Coordinated Traffic Signals: A Deep Learning Approach. Technical Report. Urban Mobility & Equity Center.
- Fletcher, R., 1971. A general quadratic programming algorithm. IMA Journal of Applied Mathematics 7, 76–91.
- Hao, P., Wu, G., Boriboonsomsin, K., Barth, M.J., 2018. Eco-approach and departure (ead) application for actuated signals in real-world traffic. IEEE Transactions on Intelligent Transportation Systems 20, 30–40.
- Heidari, M., Chow, V.T., Kokotović, P.V., Meredith, D.D., 1971. Discrete differential dynamic programming approach to water resources systems optimization. Water Resources Research 7, 273–282.
- Hounsell, N., McDonald, M., 2001. Urban network traffic control. Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering 215, 325–334.
- Jacobson, D.H., Mayne, D.Q., 1970. Differential dynamic programming. 24, Elsevier Publishing Company.
- Jamshidinejad, A., Papamichail, I., Papageorgiou, M., De Schutter, B., 2017. Sustainable model-predictive control in urban traffic networks: Efficient solution based on general smoothening methods. IEEE Transactions on Control Systems Technology 26, 813–827.
- Katsaros, K., Kernchen, R., Dianati, M., Rieck, D., 2011. Performance study of a green light optimized speed advisory (glosa) application using an integrated cooperative its simulation platform, in: 2011 7th International Wireless Communications and Mobile Computing Conference, IEEE. pp. 918–923.
- Kosmatopoulos, E., Papageorgiou, M., Bielefeldt, C., Dinopolou, V., Morris, R., Mueck, J., Richards, A., Weichenmeier, F., 2006. International comparative field evaluation of a traffic-responsive signal control strategy in three cities. Transportation Research Part A: Policy and Practice 40, 399–413.
- Koukoumidis, E., Peh, L.S., Martonosi, M.R., 2011. Signalguru: leveraging mobile phones for collaborative traffic signal schedule advisory, in: Proceedings of the 9th international conference on Mobile systems, applications, and services, pp. 127–140.
- Lawitzky, A., Wollherr, D., Buss, M., 2013. Energy optimal control to approach traffic lights, in: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE. pp. 4382–4387.
- Lu, M., Ferragut, J., Kutilla, M., Chen, T., 2020. Next-generation wireless networks for v2x, in: 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), IEEE. pp. 1–5.
- Ma, J., Zhou, F., Huang, Z., Melson, C.L., James, R., Zhang, X., 2018. Hardware-in-the-loop testing of connected and automated vehicle applications: a use case for queue-aware signalized intersection approach and departure. Transportation Research Record 2672, 36–46.
- Mahler, G., Vahidi, A., 2012. Reducing idling at red lights based on probabilistic prediction of traffic signal timings, in: 2012 American Control Conference (ACC), IEEE. pp. 6557–6562.
- Murray, D., Yakowitz, S., 1984. Differential dynamic programming and newton's method for discrete optimal control problems. Journal of Optimization Theory and Applications 43, 395–414.
- Murray, D.M., Yakowitz, S.J., 1979. Constrained differential dynamic programming and its application to multireservoir control. Water Resources Research 15, 1017–1027.
- Papageorgiou, M., Diakaki, C., Dinopolou, V., Kotsialos, A., Wang, Y., 2003. Review of road traffic control strategies. Proceedings of the IEEE 91, 2043–2067.
- Papageorgiou, M., Leibold, M., Buss, M., 2015. Optimierung. volume 4. Springer.
- Sanchez, M., Cano, J.C., Kim, D., 2006. Predicting traffic lights to improve urban traffic fuel consumption, in: 2006 6th International Conference on ITS Telecommunications, IEEE. pp. 331–336.
- Scheegans, L.E., Heckmann, K., Hoyer, R., 2022. Exploiting stage information for prediction of switching times of traffic actuated signals using machine learning, in: 2022 12th International Conference on Advanced Computer Information Technologies (ACIT), IEEE. pp. 544–548.
- Stahlmann, R., Möller, M., Brauer, A., German, R., Eckhoff, D., 2016. Technical evaluation of glosa systems and results from the field, in: 2016 IEEE Vehicular Networking Conference (VNC), IEEE. pp. 1–8.
- Sun, C., Guanetti, J., Borrelli, F., Moura, S.J., 2020. Optimal eco-driving control of connected and autonomous vehicles through signalized intersections. IEEE Internet of Things Journal 7, 3759–3773.
- Tassa, Y., Mansard, N., Todorov, E., 2014. Control-limited differential dynamic programming, in: 2014 IEEE International Conference on Robotics and Automation (ICRA), IEEE. pp. 1168–1175.
- Typaldos, P., Kalogianni, I., Mountakis, K.S., Papamichail, I., Papageorgiou, M., 2020a. Vehicle trajectory specification in presence of traffic lights with known or uncertain switching times. Transportation research record 2674, 53–66.
- Typaldos, P., Papamichail, I., Papageorgiou, M., 2020b. Minimization of fuel consumption for vehicle trajectories. IEEE Transactions on Intelligent Transportation Systems 21, 1716–1727.
- Typaldos, P., Volakakis, V., Papageorgiou, M., Papamichail, I., 2021. Vehicle-based trajectory specification in presence of traffic lights with stochastic switching times. IFAC-PapersOnLine 54, 298–305.
- Xie, Z., Liu, C.K., Hauser, K., 2017. Differential dynamic programming with nonlinear constraints, in: 2017 IEEE International Conference on Robotics and Automation (ICRA), IEEE. pp. 695–702.
- Yakowitz, S., 1986. The stagewise kuhn-tucker condition and differential dynamic programming. IEEE transactions on automatic control 31, 25–30.
- Yang, H., Almutairi, F., Rakha, H., 2020. Eco-driving at signalized intersections: A multiple signal optimization approach. IEEE Transactions on Intelligent Transportation Systems 22, 2943–2955.
- Yang, H., Rakha, H., Ala, M.V., 2016. Eco-cooperative adaptive cruise control at signalized intersections considering queue effects. IEEE Transactions on Intelligent Transportation Systems 18, 1575–1585.
- Zhao, W., Ngoduy, D., Shepherd, S., Liu, R., Papageorgiou, M., 2018. A platoon based cooperative eco-driving model for mixed automated and human-driven vehicles at a signalised intersection. Transportation Research Part C: Emerging Technologies 95, 802–821.