Experimental Validation of a semi-distributed SQP method for optimal coordination of automated vehicles at intersections

Robert Hult*[1] | Mario Zanon[2] | Gianluca Frison[3] | Sébastien Gros[1] | Paolo Falcone[1]

[1]Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden
[2]IMT School of Advanced Studies, Lucca, Italy
[3]Department of Microsystems Engineering, University of Freiburg, Freiburg, Germany

Correspondence
*Corresponding author. Email: robert.hult@chalmers.se

Present Address
Department of Electrical Engineering, Chalmers University of Technology, Hörsalsvägen 11, SE-41296 Göteborg, Sweden

Summary

In this paper, we study the optimal coordination of automated vehicles at intersections. The problem can be stated as an optimal control problem, which can be decomposed into one NLP which schedules access to the intersection and one optimal control problem per vehicle. The decomposition enables a bi-level MPC, where an outer control loop schedules access to the intersection, and inner control loops compute the appropriate vehicle commands. We discuss a practical implementation of the bi-level controller where the NLP is solved with a tailored SQP algorithm which enables distribution of most computation to the vehicles. Results from an extensive experimental campaign are presented, where the bi-level controller and the semi-distributed SQP is implemented on a test setup consisting of three automated vehicles. In particular, we show that the vehicle-level controller can enforce the scheduled intersection access within a relevant accuracy, and that the bi-level controller can handle large perturbations and large communication delays, which makes the scheme applicable in practical scenarios. Finally, the use of wireless communication introduce delays in the outer control loop and to allow faster feedback, we introduce a Real Time Iteration (RTI) like variation of the bi-level controller. Experimental and simulated results indicate that the RTI-variation offer comparable performance using less computation and communication.

KEYWORDS:
Automated Vehicles, Intersection Coordination, Distributed Model Predictive Control, Distributed Nonlinear programming

1 | INTRODUCTION

The current trend towards automation of road vehicles can be expected to continue, and eventually virtually all vehicles will be fully automated and communicating. This technology can be leveraged to obtain synergistic effects through cooperation between the automated vehicles, and thereby enable drastic improvements to the traffic system. In this paper we discuss an algorithm necessary for one such improvement: the automation of intersection crossings. With all vehicles automated, communicating and cooperative, the traffic-lights, signs and rules used today could be removed and the vehicles could instead rely on automated coordination controllers. As discussed in [1], the potential benefits include increased safety, increased energy efficiency and higher traffic throughput.

However, there are several challenges that must be addressed before coordination controllers can be applied in practice. Most importantly, such controllers must be able to guarantee that no collisions occur, and in particular, the guarantees must be applicable to scenarios with uncertainty. This includes handling of unexpected events and the on-line re-coordination of vehicles in the presence of new information. Furthermore, a useful coordination algorithm must be scalable to be relevant for more than small scenarios. However, since finding the optimal collision free motion profiles for vehicles crossing an intersection is a combinatorial problem, there are computational scalability issues. In fact, determining the existence of even one collision free solution has been shown to be an NP-hard problem in the general case[2]. Moreover, it is a known problem that vehicle-to-vehicle (V2V) communication systems have capacity limits, so that the number of vehicles that can communicate simultaneously is restricted[3]. A practically useful coordination algorithm must therefore also scale well in terms of both how often communication is required and the data volumes involved.

Although the application of intersection coordination controllers lies in the future, a number of contributions have been made during the last decade, most of which are surveyed in[4] and[5]. In much of the current literature, various heuristics are used to address the challenges of the problem. In these works, the motion profile is typically the result of a rule-based controller which switches between discrete behavioral modes[6,7,8,9], or is obtained from a restricted space, e.g. trapezoidal[10] or linear[11] velocity profiles. On the other hand, a number of approaches based on Optimal Control (OC) formulations of the coordination problem can be found in the literature, e.g.[12,13,14,15,16,17,18,19,20,21,?] and[?]. In most cases, the computation of a crossing order is separated from the computation of the optimal state and control trajectories to avoid the combinatorial explosion of the solution space. For instance, in[14,16,19] variations of "First-Come-First-Served" policies are used to produce a crossing order and OC problems constrained to satisfy this order are solved afterwards. Using a similar strategy,[15] leverages results from polling-systems to compute the crossing order, while in[22] and[?] mixed-integer quadratic programing is used to compute an approximately optimal crossing order. A different approach is taken in[12], where a heuristic gives a decision order rather than a crossing order. The vehicles thereafter sequentially solve optimal control problems, where each vehicle is constrained to avoid collisions with the vehicles that precedes it in the decision order. The application of OC formulations to closed-loop control is considered in e.g.[13,21] and[23].

As discussed in[1], the benefits of OC approaches in general include their ability to consider a wider range of applicable motion profiles, include constraints and use explicitly stated objective functions. Contrary to many heuristic methods, explicit objective functions gives an additional degree of control over the characteristics of a coordinated solution, which is a desirable property. Given the severity of collisions, closed-loop control, i.e. the recalculation of control commands based on measurements of the system state, is a necessity to handle the uncertainty that is present in real scenarios. In many cases, OC schemes can leverage well established theory to derive properties of closed-loop control schemes, account for various forms of uncertainty and to construct efficient solution algorithms.

In this paper we use the OC formulation of the coordination problem first presented in[22] but focus on finding the optimal solution for a given crossing order. With this formulation, the problem is given a hierarchical structure, where optimal, collision free intersection occupancy time-slots are obtained as the solution to a Nonlinear Program (NLP), and the optimal state and control trajectories as the solution to optimal control problems that are separable between the vehicles. The structure enables a bi-level Model Predictive Control (MPC) architecture where coordination is separated from vehicle control. In particular, the upper, *intersection*-level, controller computes and updates optimal, non-overlapping time-slots based on the current vehicle states, and the lower, *vehicle*-level, controllers compute the optimal control commands for the vehicles, given a time-slot and current state.

In earlier work, we proposed a semi-distributed Sequential Quadratic Programing (SQP) for the solution of the timeslot NLP[24]. The algorithm was extended in[25] where a convergence proof was also given. We established the persistent feasibility of the bi-level MPC scheme and discussed robustness aspects in[23], and presented experimental results. Extensions to Economic Nonlinear MPC were presented in[?], and a comparison of[24] and[25] was given in[?], supported by experimental data.

The SQP procedure in[24] and[25] has the property that most computations can be performed on-board the involved vehicles and the algorithm's internal information passing can be performed using V2V communication. This enables a practical scenario where the SQP procedure is used to close the intersection-level control loop (repeatedly solve the timeslot NLP on-line) in a semi-distributed manner. While such a scheme has several desirable properties, it is necessary to evaluate its practical usefulness in a real world setting. First, the algorithmic performance needs to be assessed for real scenarios. In particular, the effects of delays inherent to practical communication systems must be studied. Second, the effects of both poor algorithmic performance and real-world perturbations on the performance of the bi-level controller must be investigated and possible issues addressed.

FIGURE 1 Subfigure (a) contains a schematic illustration of the scenarios considered in this paper. Subfigure (b) illustrates how the intersection is modeled: the arrows shows the fixed paths of the vehicles, and the red square illustrates the zone inside the intersection where collisions can occur.

To this end, we describe a practical implementation of the the bi-level controller in this paper, where the intersection-level control loop is closed using the SQP presented in [24], [25]. In particular, we detail the application of the bi-level controller on a test-setup consisting of three automated vehicles, where the SQP is solved in a semi-distributed fashion using V2V communication, and the vehicle-level control loop is closed using the high-performance QP solver HPMPC [26]. Furthermore, we introduce modifications to the bi-level controller which increase the practical applicability of the scheme. First, a relaxation of vehicle-level MPC problem is presented, which resolves infeasibility issues inherent to the formulation in [24], [25]. Second, two modifications to the intersection-level controller are introduced to handle the large computational delays that can arise due to the execution of the SQPs over a wireless network. In particular, we propose a scheme where the intersection control loop is closed in a Real Time Iteration (RTI) fashion [27] to allow faster feedback and less communication. That is, instead of solving the NLP to convergence, the intersection-level control loop consists of performing one SQP iteration and thereafter applying the resulting timeslot. Moreover, we present results from an extensive experimental campaign where the implementation was evaluated. We discuss the algorithmic performance and provide a detailed study on the SQP execution times, where the experimental data is compared to ideal cases. We also analyze the difference between the different modifications to the intersection-level controllers. Comparative data is provided for both experimental and simulated cases, where the system is subjected to both large and small perturbations.

The remainder of the paper is organized as follows: In Section 2 we introduce the problem and state the optimal control problem. In Section 3 we detail the semi-distributed SQP method. In Section 4 we detail a practical implementation of the SQP and bi-level control. In Section 5 we present and discuss experimental results. The paper is concluded in Section 6.

2 | PROBLEM FORMULATION

In this paper we consider problems such as that illustrated in Figure 1 a, where N_a vehicles approach an intersection equipped with a central coordinating unit. We assume that all involved vehicles are automated, cooperative and participate in the coordination procedure and that no other, non-cooperative entities (e.g., pedestrians and bicyclists) are present. For simplicity, we also assume that no vehicles makes turns or change lanes.

2.1 | Modeling

We assume that the vehicles move along predefined paths and that the vehicle dynamics along the paths can be described by

$$\dot{x}_i(t) = A_i^c x_i(t) + B_i^c u_i(t), \tag{1}$$

where $x_i(t) \in \mathbb{R}^n$, $u_i(t) \in \mathbb{R}^m$ and $A_i^c \in \mathbb{R}^{n \times n}$, $B_i^c \in \mathbb{R}^{n \times m}$. Specifically, the state vector is such that $x_i(t) = (p_i(t), y_i(t))$, where $p_i(t) \in \mathbb{R}$ is the position of the center of the vehicle on its path and $y_i(t) \in \mathbb{R}^{n-1}$ collects all non position states (e.g., velocity

and acceleration). Moreover, the vehicle state- and control trajectories are subject to constraints of the form

$$D_i x_i(t) + G_i u_i(t) \geq b_i, \tag{2}$$

capturing e.g. actuation limitations and passenger comfort restrictions.

We consider only D_i, G_i, A_i^c, B_i^c such that $\dot{p}(t) \geq 0$, i.e. the dynamics and constraints are such that no vehicle can reverse.

As illustrated in Figure 1 b, we define the intersection as an interval $[p_i^{in}, p_i^{out}]$ on the path of each vehicle such that collisions are avoided if $p_i(t) \in [p_i^{in}, p_i^{out}] \Rightarrow p_j(t) \notin [p_j^{in}, p_j^{out}]$ hold, for all vehicles $i \neq j$[1]. Furthermore, we define the time-of-entry, t_i^{in}, and time-of-clearance, t_i^{out}, of the intersection through

$$p_i(t_i^{in}) = p_i^{in} \text{ and } p_i(t_i^{out}) = p_i^{out}, \tag{3}$$

respectively. Collision avoidance is thereby ensured if

$$t_i^{out} \leq t_j^{in} \tag{4}$$

for all vehicle pairs (i, j) such that vehicle i crosses the intersection before vehicle j. In the remainder of the paper we denote $T_i = (t_i^{in}, t_i^{out})$ the *time-slot* of vehicle i, and say that a vehicle conforms to a time-slot if it only occupies the intersection within $[t_i^{in}, t_i^{out}]$.

2.2 | Optimal Control formulation

With the objective

$$J_i(x_i(t), u_i(t)) = V_{i,f}(x_i(t_f)) + \int_0^{t_f} \ell_i(x_i(t), u_i(t)) dt, \tag{5}$$

where $V_{i,f}(x_i(t_f))$ and $\ell_i(x_i(t), u_i(t))$ are convex and quadratic, and t_f is fixed, we state problem of optimally coordinating the vehicles through the intersection as

$$\min_{T, x(t), u(t)} \sum_{i=1}^{N_a} J_i(x_i(t), u_i(t)) \tag{6a}$$

$$\text{s.t.} \quad x_i(0) = \hat{x}_{i,0} \qquad\qquad i \in \mathbb{I}_{[1,N_a]}, \tag{6b}$$

$$\dot{x}_i(t) = A_i^c x_i(t) + B_i^c u_i(t), \qquad\qquad i \in \mathbb{I}_{[1,N_a]}, \tag{6c}$$

$$D_i x_i(t) + G_i u_i(t) \geq b_i,, \qquad\qquad i \in \mathbb{I}_{[1,N_a]}, \tag{6d}$$

$$p_i(t_i^{in}) = p_i^{in}, \quad p_i(t_i^{out}) = p_i^{out}, \qquad\qquad i \in \mathbb{I}_{[1,N_a]}, \tag{6e}$$

$$t_i^{out} \leq t_{i+1}^{in} \qquad\qquad i \in \mathbb{I}_{[1,N_a-1]}. \tag{6f}$$

Here, $\mathbb{I}_{[a,b]} = \{a, \dots, b\}$ for integers $a < b$, $x(t) = (x_1(t), \dots, x_{N_a}(t))$, $u(t) = (u_1(t), \dots, u_{N_a}(t))$, $T = (T_1, \dots, T_{N_a})$ and $\hat{x}_{i,0}$ is the initial state of vehicle i, and the vehicles are ordered such that vehicle i crosses before vehicle $i+1$.

2.2.1 | Decomposition

It was shown in[22] that the coordination problem can be decomposed in a hierarchical fashion, where the time-slot schedule T is the solution of a nonlinear program (NLP), and the vehicle state and control trajectories $x_i(t), u_i(t)$ are the solution to separable vehicle optimal control problems.

With this decomposition, the following NLP give the optimal time-slot schedule T for a given order S

$$\min_T \sum_{i=1}^{N_a} V_i(\hat{x}_{i,0}, T_i) \tag{7a}$$

$$\text{s.t.} \quad T_i \in \text{domain}(V_i(\hat{x}_{i,0}, T_i)), \qquad\qquad i \in \mathbb{I}_{[1,N_a]}, \tag{7b}$$

$$t_i^{out} \leq t_{i+1}^{in} \qquad\qquad i \in \mathbb{I}_{[1,N_a]}, \tag{7c}$$

---

[1]We want to emphasize that the definition of the intersection easily can be subdivided into several mutual-exclusion zones, each with its own start and stop position. However, for simplicity of presentation the paper is developed with the single zone shown in Figure 1 b.

where $V_i(\hat{x}_{i,0}, T_i)$ is defined as the optimal value function of the vehicle optimal control problem

$$V_i(\hat{x}_{i,0}, T_i) = \min_{x_i(t), u_i(t)} \quad J_i(x_i(t), u_i(t)) \quad \text{s.t.} \quad x_i(0) = \hat{x}_{i,0}, \ (1), \ (2), \ p_i(t_i^{\text{in}}) = p_i^{\text{in}}, \ p_i(t_i^{\text{out}}) = p_i^{\text{out}}. \tag{8}$$

For the optimal T_i, (8) gives the optimal state and control trajectories $x_i(t), u_i(t)$.

2.2.2 | Discretization

For practical reasons, we consider piece-wise constant inputs and discretization of the vehicle dynamics using the sampling time T_s when we solve the vehicle problem (8). More precisely, we define $x_{i,k} = x_i(t_k)$ and $u(t) = u_{i,k}, \forall t \in [t_k, t_{k+1}[$, with $t_k = kt_s$, and the state update function $x_{i,k+1} = A_i x_{i,k} + B_i u_{i,k}$, where $A_i = \mathcal{A}_i(t_s) = \exp(A_i^c t_s)$, $B_i = \mathcal{B}_i(t_s) = \int_0^{t_s} \exp(A_i^c(t_s - s))B_i^c ds$ and where $x_{i,k} = (p_{i,k}, v_{i,k})$.

Since the discrete position is defined only at t_k, (3) only defines values of t_i^{in} and t_i^{out} that are integer multiples of t_s in the discrete time case. To allow t_i^{in} and t_i^{out} to assume continuous values, we define a continuous time representation of the position using the discrete time state- and control sequences as

$$p_i^d(t, w_i) = [1, \mathbf{0}^{n-1}] \left(\mathcal{A}_i(t - t_k)x_{i,k} + \mathcal{B}_i(t - t_k)u_{i,k} \right), \quad k = \text{floor}(t/t_s), \tag{9}$$

where $w_i = (x_{i,0}, u_{i,0}, \dots, x_{i,N-1}, u_{i,N-1}, x_{i,N})$ and where $t_f = Nt_s$. The discrete time statement of the objective function is

$$J_i^d(w_i) = x_{i,N}^\top P_i x_{i,N} + \sum_{k=0}^{N-1} \frac{1}{2} \begin{bmatrix} x_{i,k} \\ u_{i,k} \end{bmatrix}^\top Q_i \begin{bmatrix} x_{i,k} \\ u_{i,k} \end{bmatrix} + q_i^\top \begin{bmatrix} x_{i,k} \\ u_{i,k} \end{bmatrix} \tag{10}$$

whereby we have the discrete time formulation of (8) as

$$V_i(\hat{x}_{i,0}, T_i) = \min_{w_i} \quad J_i^d(w_i) \tag{11a}$$

$$\text{s.t.} \quad x_{i,0} = \hat{x}_{i,0} \tag{11b}$$

$$x_{i,k+1} = A_i x_{i,k} + B_i u_{i,k}, \qquad k \in \mathbb{I}_{[0,N-1]}, \tag{11c}$$

$$D_i x_{i,k} + G_i u_{i,k}(t) \geq b_i,, \qquad k \in \mathbb{I}_{[0,N-1]}, \tag{11d}$$

$$p_i^d(t_i^{\text{in}}, w_i) = p_i^{\text{in}}, \tag{11e}$$

$$p_i^d(t_i^{\text{out}}, w_i) = p_i^{\text{out}}, \tag{11f}$$

which can be solved for real-valued t_i^{in} and t_i^{out} in $[0, Nt_s]$. We note here that for w_i which satisfies (11c), $p_i^d(t, w_i)$ is K-times continuously differentiable, where K is the relative degree of (1) with the position $p_i(t)$ as the output[24].

2.2.3 | Problem Properties

The constraint set domain$(V_i(\hat{x}_{i,0}, T_i))$ in (7b) is implicitly defined as the set of T_i for which the optimization problem (11) is feasible given the initial state $\hat{x}_{i,0}$. However, it was shown in[22] that domain$(V_i(\hat{x}_{i,0}, T_i))$ can be written as

$$h_i(\hat{x}_{i,0}, T_i) = \begin{bmatrix} L^{\text{in}}(\hat{x}_{i,0}) - t_i^{\text{in}} \\ t_i^{\text{in}} - U^{\text{in}}(\hat{x}_{i,0}) \\ U(\hat{x}_{i,0}, t_i^{\text{in}}) - t_i^{\text{out}} \\ t_i^{\text{out}} - L(\hat{x}_{i,0}, t_i^{\text{in}}) \end{bmatrix} \geq 0, \tag{12}$$

where $L^{\text{in}}(\hat{x}_{i,0}), U^{\text{in}}(\hat{x}_{i,0}), U(\hat{x}_{i,0}, t_i^{\text{in}})$ and $L(\hat{x}_{i,0}, t_i^{\text{in}})$ are defined as the solutions to the NLPs

$$L^{\text{in}}(\hat{x}_{i,0}) = \min_{w_i, t} t \quad \text{s.t.} \quad (11b), (11c), (11d), p_i^d(t, w_i) = p_i^{\text{in}}, \tag{13}$$

$$U^{\text{in}}(\hat{x}_{i,0}) = \max_{w_i, t} t \quad \text{s.t.} \quad (11b), (11c), (11d), p_i^d(t, w_i) = p_i^{\text{in}}, \tag{14}$$

$$L(\hat{x}_{i,0}, t_i^{\text{in}}) = \min_{w_i, t} t \quad \text{s.t.} \quad (11b), (11c), (11d), p_i^d(t_i^{\text{in}}, w_i) = p_i^{\text{in}}, p_i^d(t, w_i) = p_i^{\text{out}}, \tag{15}$$

$$U(\hat{x}_{i,0}, t_i^{\text{in}}) = \max_{w_i, t} t \quad \text{s.t.} \quad (11b), (11c), (11d), p_i^d(t_i^{\text{in}}, w_i) = p_i^{\text{in}}, p_i^d(t, w_i) = p_i^{\text{out}}. \tag{16}$$

That is, t_i^{in} must lie between the earliest and latest time-of-entry that the vehicle can perform. Similarly, for a specified time-of-entry t_i^{in}, t_i^{out} must lie between the earliest and latest time-of-clearance the vehicle can perform. Moreover, it was shown in[24]

that if a mild technical assumption holds, the optimizers of the linear programs (LPs)

$$w_i^{\mathrm{ub}}(\hat{x}_{i,0}, t_i^{\mathrm{in}}) = \arg\min_{w_i} p_{i,N} \quad \text{s.t.} \quad (11b), (11c), (11d), \ p_i^d(t_i^{\mathrm{in}}, w_i) = p_i^{\mathrm{in}}, \tag{17}$$

$$w_i^{\mathrm{lb}}(\hat{x}_{i,0}, t_i^{\mathrm{in}}) = \arg\max_{w_i} p_{i,N} \quad \text{s.t.} \quad (11b), (11c), (11d), \ p_i^d(t_i^{\mathrm{in}}, w_i) = p_i^{\mathrm{in}}. \tag{18}$$

are also optimizers of (15) and (16) respectively. Consequently, (15) and (16) can be evaluated by first solving the LPs (17), (18) and thereafter solving

$$p_i^d(U(\hat{x}_{i,0},, t_i^{\mathrm{in}}), w_i^{\mathrm{ub}}(\hat{x}_{i,0}, t_i^{\mathrm{in}})) - p_i^{\mathrm{out}} = 0 \tag{19}$$

$$p_i^d(L(\hat{x}_{i,0},, t_i^{\mathrm{in}}), w_i^{\mathrm{lb}}(\hat{x}_{i,0}, t_i^{\mathrm{in}})) - p_i^{\mathrm{out}} = 0 \tag{20}$$

for $U(\hat{x}_{i,0}, t_i^{\mathrm{in}})$ and $L(\hat{x}_{i,0}, t_i^{\mathrm{in}})$. The bounds $L^{\mathrm{in}}(\hat{x}_{i,0})$ and $U^{\mathrm{in}}(\hat{x}_{i,0})$ can be obtained similarly. In the remainder of the paper we will for notational convenience only include the explicit dependence on $\hat{x}_{i,0}$ in $h_i(\hat{x}_{i,0}, T_i)$ and $V_i(\hat{x}_{i,0}, T_i)$ when necessary.

2.3 | Receding horizon implementation

In order to reject perturbations and compensate for model inaccuracies, the solution to the optimal coordination problem can be applied in a receding horizon fashion in a model predictive controller (MPC). In particular, the decomposed formulation offers a natural separation between coordination, i.e. resolving potential collisions, and vehicle control. This enables a bi-level control structure: an outer control loop computes collision free time-slots by solving (7) at the current state, while inner control loops compute the vehicle control command $u_{i,0}$ for a given time-slot T_i at the current state by solving (11). Feedback is thereby due to updated estimates of the current state in the inner control loops, which translates into modifications of the cost function $V_i(\hat{x}_{i,0}, T_i)$ and the constraint set $h_i(\hat{x}_{i,0}, T_i)$ in the outer control loop. The bi-level scheme has the benefit that perturbations that arise for one vehicle are counteracted by all vehicles, i.e the optimal timeslot of vehicle i at time k, $T_i(X_k)$, is a function of the state of all vehicles $X_k = (x_{1,k}, \ldots, x_{N_a,k})$ through (7).

Note that in principle, the time-slot schedule could be computed once and the rejection of possible perturbations could be left to the inner control loops. However, by closing the outer control loop the system can (a) reject larger perturbations by adjusting the time-slot schedule T and thereby provide collision avoidance in more demanding scenarios and (b) continuously improve the solution. The properties of the bi-level controller are discussed in [23], where the closed-loop system is shown to be persistently feasible and stable in the nominal case.

3 | A SEMI-DISTRIBUTED SQP METHOD

While the computations in the vehicle-level control loop can be performed independently by all vehicles, the upper control loop relies on the solution of the NLP (7) and information from all vehicles. Considering the intended application, a solution algorithm in which much of the computations can be parallelized and performed on board the vehicles is desirable as it improves scalability. However, if computations are performed on board the vehicles, the algorithm must be executed with information exchanged over the V2V communication network. As is reported in e.g. [3], there are scalability issues with the current V2V technology and a solution algorithm should therefore not require frequent nor large data exchange. Consequently, second-order optimization methods are preferable to first-order ones, as the former in general needs fewer iterations to find a solution to the problem. Therefore, a semi-distributed sequential programming (SQP) algorithm was proposed in [24], [25]. In the remainder of the section we detail the semi-distributed SQP.

3.1 | SQP

Using the developments of Section 2.2.3, we rewrite NLP (7) as

$$\min_{T} \ V(T) \quad \text{s.t.} \ h(T) \geq 0, \tag{21}$$

where $V(T) = \sum_{i=1}^{N_a} V_i(T_i)$ and we lumped constraints (7b)-(7c) in the function $h(T)$. The associated Lagrange function is defined as $\mathcal{L}(T, \mu) := V(T) - \mu^\top h(T)$, where $\mu = (\mu_1, \ldots, \mu_{N_a}, \mu_s)$. Here, μ_i are the Lagrange multipliers of the constraint $h_i(T_i) \geq 0$ and μ_s the multipliers of the precedence constraints (7c), which we write as $h_s(T) \geq 0$

Starting from an initial guess $z^{(0)} = (\boldsymbol{T}^{(0)}, \mu^{(0)})$, SQP iteratively updates the primal-dual solution candidate $z^{(c)}$ using

$$z^{(c+1)} = z^{(c)} + \alpha^{(c)} \Delta z^{(c)}, \tag{22}$$

with $\alpha^{(c)} \in (0, 1]$ and $\Delta z^{(c)} = (\Delta \boldsymbol{T}^{(c)}, \tilde{\mu}^{(c)} - \mu^{(c)})$. Here, $(\Delta \boldsymbol{T}^{(c)}, \tilde{\mu}^{(c)})$ is the primal-dual solution of the Quadratic Programming (QP) subproblem

$$\min_{\Delta \boldsymbol{T}} \quad \frac{1}{2} \Delta \boldsymbol{T}^\top H^{(c)} \Delta \boldsymbol{T} + \nabla_{\boldsymbol{T}} V(\boldsymbol{T}^{(c)})^\top \Delta \boldsymbol{T} \tag{23a}$$

$$\text{s.t.} \quad h(\boldsymbol{T}^{(c)}) + \nabla_{\boldsymbol{T}} h(\boldsymbol{T}^{(c)})^\top \Delta \boldsymbol{T} \geq 0, \tag{23b}$$

where $H^{(c)}$ is a positive-definite approximation of the Lagrange function Hessian $\nabla_{\boldsymbol{T}}^2 \mathcal{L}(\boldsymbol{T}^{(c)}, \mu^{(c)})$. Variants of SQP differ primarily in the computations of the step size $\alpha^{(k)}$ and the Lagrange function Hessian approximation $H^{(k)}$, and we describe next the methods employed to solve (7). For more details on SQP see e.g.,[28].

3.1.1 | Hessian Approximation

To ensure that the QP subproblems (23) are convex, it is in general required that the reduced Hessian is positive-definite. While there are several ways of enforcing positive-definiteness of the reduced Hessian, we adopt the strategy of adding enough curvature in all negative-curvature directions to ensure that the full Hessian is positive-definite. In particular, we note that NLP (21) is such that the Hessian has the block-diagonal form

$$\nabla_{\boldsymbol{T}}^2 \mathcal{L}(\boldsymbol{T}^{(c)}, \mu^{(c)}) = \text{diag}(L_i(T_1^{(c)}, \mu_1^{(c)}), \dots, L_{N_a}(T_{N_a}^{(c)}, \mu_{N_a}^{(c)})) \tag{24}$$

where $L_i(T_i^{(c)}, \mu_i^{(c)}) = \nabla_{T_i}^2 V_i(T_i^{(c)}) + \langle \mu_i^{(c)}, \nabla_{T_i}^2 h_i(T_i^{(c)}) \rangle$. We define a positive-definite approximation to $\nabla_{\boldsymbol{T}}^2 \mathcal{L}(\boldsymbol{T}^{(c)}, \mu^{(c)})$ as $H^{(c)} = \text{diag}(H_1^{(c)}, \dots, H_{N_a}^{(c)})$, where $H_i^{(c)} = E_i^{(c)} D_i^{(c)} E_i^{{(c)}^\top}$. Here, $D_i^{(c)}$ is a diagonal matrix where the j^{th} diagonal element is $d_{i,j}^{(c)} = \max(e_{i,j}^{(c)}, \varepsilon)$, where $e_{i,j}$ is the j^{th} eigenvalue of $L_i(T_i^{(c)}, \mu_i^{(c)})$ and $\varepsilon > 0$ is a constant. The columns of $E_i^{(c)}$ are the normalized eigenvectors corresponding to the eigenvalues in D_i. Note that, due to the small size of the blocks $L_i(T_i^{(c)}, \mu_i^{(c)})$, the required eigenvalue decomposition is very cheap to compute.

3.1.2 | Step size selection

In order to guarantee convergence of SQP algorithms, the step-size $\alpha^{(c)}$ must be selected such that progress towards a solution to the problem is made. In this paper, we employ a line search on the so called ℓ_1 merit function, which is defined as

$$M(\boldsymbol{T}^{(c)}) = V(\boldsymbol{T}^{(c)}) + \rho^{(c)} ||h^-(\boldsymbol{T}^{(c)})||_1, \tag{25}$$

where $h^-(\boldsymbol{T}^{(c)}) = \min(h(\boldsymbol{T}^{(c)}), 0)$ and $\rho^{(c)}$ is a parameter chosen so that $\rho^{(c)} > ||\mu^{(c)}||_1$. Progress towards a solution is ensured when $\alpha^{(c)}$ is selected such that the Armijo condition is satisfied:

$$M(\boldsymbol{T}^{(c)} + \alpha^{(c)} \Delta \boldsymbol{T}^{(c)}) \leq M(\boldsymbol{T}^{(c)}) + \gamma D_{\Delta \boldsymbol{T}^{(c)}} M(\boldsymbol{T}^{(c)}) \alpha^{(c)}, \tag{26}$$

where $\gamma \in (0, 0.5]$ and $D_{\Delta \boldsymbol{T}^{(c)}} M(\boldsymbol{T}^{(c)})$ is the derivative of $M(\boldsymbol{T})$ in the direction of $\Delta \boldsymbol{T}^{(c)}$, evaluated at $\boldsymbol{T}^{(c)}$. Provided that $\Delta \boldsymbol{T}^{(c)}$ is a descent direction on (25), $\alpha^{(c)}$ which satisfies (26) exists and can be found by so-called backtracking, i.e. by successively decreasing $\alpha^{(c)}$ from 1 until (26) is satisfied[28].

Since the constraint set $h_i(T_i) \geq 0$ defines the set of feasible parameters for the parametric QP (11), and $V_i(T_i)$ is the optimal value function for the same QP, we note that $V_i(T_i)$, and thereby $M(\boldsymbol{T})$, is undefined when $h_i(T_i) \ngeq 0$. We resolve this issue by using the projection based method presented in[25], where the merit function is evaluated at a projected point $\mathcal{P}(T_i^{(c)} + \alpha^{(c)} \Delta T_i^{(c)})$ rather than at $T_i^{(c)} + \alpha^{(c)} \Delta T_i^{(c)}$. Here, the projection operator is defined as

$$\mathcal{P}(T_i) := \left(t_i^{\text{in}}, \ \min(U(t_i^{\text{in}}), \max(L(t_i^{\text{in}}), t_i^{\text{out}})) \right). \tag{27}$$

It was shown in[25] that if $\Delta \boldsymbol{T}^{(c)}$ is a descent direction on (25), small enough $\alpha^{(c)}$ exists which satisfies (26), giving that one can backtrack on $M(\mathcal{P}(\boldsymbol{T}^{(c)} + \alpha^{(c)} \Delta \boldsymbol{T}^{(c)}))$ as well. Note that with this modification, the ℓ_1 merit function reads $M(\mathcal{P}(\boldsymbol{T})) = V(\mathcal{P}(\boldsymbol{T})) + \rho ||h_s^-(\mathcal{P}(\boldsymbol{T})) \lambda \mathbb{Z}||$ and the primal-dual update is

$$T_i^{(c+1)} = \mathcal{P}\left(T_i^{(c)} + \alpha^{(c)} \Delta T_i^{(c)} \right), \quad i \in \mathbb{I}_{[1, N_a]} \tag{28a}$$

$$\mu^{(c+1)} = \mu^{(c)} + \alpha^{(c)} \Delta \mu^{(c)}. \tag{28b}$$

We note that an alternative solution to the issue of non-defined $M(T)$ is to soften the position constraints (11e) and (11f) with an ℓ_1 penalty as suggested in[24]. However, in doing so the objective in the quadratic subproblem (42) will be dominated by the penalty term whenever it is evaluated at $h_i(T_i) \not\geq 0$ for one vehicle, and a sharp non-smoothness will appear at the point where the problem becomes feasible. The algorithmic performance of this method has been found to be worse than the method based on the projection (27)[?].

3.2 | Calculation of derivatives

The first and second-order derivatives of the objective function components $V_i(T_i^{(c)})$ and constraint components $h_i(T_i^{(c)})$ are required to form the QP-subproblem (23). Since $V_i(T_i^{(c)})$ is the optimal value function of the QP (11) and $h_i(T_i^{(c)})$ is evaluated using the solution to the LPs (18), (17), the derivatives are evaluated using tools from parametric sensitivity analysis. In particular, we have that

$$\frac{\mathrm{d}V_i(T_i)}{\mathrm{d}t_i^{\mathrm{in}}} = \frac{\partial \mathcal{L}_i(w_i(T_i), \lambda_i(T_i), v_i(T_i))}{\partial t_i^{\mathrm{in}}} = v_i^{\mathrm{in}}(T_i) \frac{\partial p_i^d \left(t_i^{\mathrm{in}}, w_i(T_i)\right)}{\partial t_i^{\mathrm{in}}}, \tag{29}$$

where $\mathcal{L}_i(w_i, \lambda_i, v_i)$ is the Lagrange Function of the QP (11)[29]. Here, $w_i(T_i)$ is the primal solution of (11) for T_i, $\lambda_i(T_i)$ is the dual solution corresponding to the constraints (11b)-(11d), and $v_i(T_i) = (v_i^{\mathrm{in}}(T_i), v_i^{\mathrm{out}}(T_i))$ the dual solution corresponding to constraints (11e), (11f). The second-order derivatives can thereafter be obtained using the chain rule, and we have, for instance,

$$\frac{\mathrm{d}^2 V_i(T_i)}{\mathrm{d}t_i^{\mathrm{in}2}} = \frac{\mathrm{d}v_i^{\mathrm{in}}}{\mathrm{d}t_i^{\mathrm{in}}} \frac{\partial p_i^d \left(t_i^{\mathrm{in}}, w_i(T_i)\right)}{\partial t_i^{\mathrm{in}}} + v_i^{\mathrm{in}}(T_i) \left(\frac{\partial^2 p_i^d \left(t_i^{\mathrm{in}}, w_i(T_i)\right)}{\partial t_i^{\mathrm{in}2}} + \frac{\partial^2 p_i^d \left(t_i^{\mathrm{in}}, w_i(T_i)\right)}{\partial t_i^{\mathrm{in}} \partial w_i} \frac{\mathrm{d}w_i}{\mathrm{d}t_i^{\mathrm{in}}} \right). \tag{30}$$

The derivatives of the constraints are obtained similarly, as is exemplified for $U(t_i^{\mathrm{in}})$ below. We have by definition that

$$p_i^d(U(t_i^{\mathrm{in}}), w_i^{\mathrm{ub}}(t_i^{\mathrm{in}})) - p_i^{\mathrm{out}} = 0, \tag{31}$$

where $w_i^{\mathrm{ub}}(t_i^{\mathrm{in}})$ is the solution to (17) for t_i^{in}. Differentiation w.r.t. t_i^{in} then gives that

$$\frac{\mathrm{d}U(t_i^{\mathrm{in}})}{\mathrm{d}t_i^{\mathrm{in}}} = - \left(\frac{\partial p_i^d \left(t_i^{\mathrm{in}}, w_i^{\mathrm{ub}}(t_i^{\mathrm{in}})\right)}{\partial t_i^{\mathrm{in}}} \right)^{-1} \frac{\partial p_i^d \left(t_i^{\mathrm{in}}, w_i^{\mathrm{ub}}(t_i^{\mathrm{in}})\right)}{\partial w_i^{\mathrm{ub}}} \frac{\mathrm{d}w_i^{\mathrm{ub}}}{\mathrm{d}t_i^{\mathrm{in}}}. \tag{32}$$

The second-order derivative is obtained by applying the chain rule to (32), but the resulting expression is rather large and is omitted her for brevity. However, the interested reader can find it in[24], and it should be noted that it includes the term $\mathrm{d}^2 w_i / \mathrm{d}t_i^{\mathrm{in}2}$. The computation of the first and second-order derivatives of $L(t^{\mathrm{in}})$ is identical.

Parametric Sensitivity Analysis

The expressions (29), (30) and (32) rely on the first and second-order sensitivity of the primal-dual solution with respect to T_i, which acts as a problem parameter in the QP (11) and LPs (18),(17). Note that for the general parametric optimization problem with free variable x and parameter p

$$\min_x q(x) \quad \text{s.t} \quad a(x, p) \geq 0 \tag{33}$$

the KKT conditions are satisfied at the solution, given that some constraint qualification hold. Denoting the primal dual solution to (33) $z^* = (x^*, \lambda^*)$, the KKT conditions are

$$\nabla \mathcal{L}(z^*, p) = 0, \quad a(x^*, p) \geq 0, \quad \lambda^* \geq 0, \quad a(x^*, p) * \lambda^* = 0, \tag{34}$$

where $\mathcal{L}(z) = q(x, p) - \lambda^\top a(x, p)$ and $*$ denotes element-wise multiplication. In particular, denoting the set of active constraints at the solution $a_{\mathbb{A}}(x^*, p)$ and the corresponding multipliers $\lambda_{\mathbb{A}}^*$, we have that

$$r = \begin{bmatrix} \nabla \mathcal{L}(z^*, p) \\ a_{\mathbb{A}}(x^*, p) \\ a_{\mathbb{A}}(x^*, p) * \lambda_{\mathbb{A}}^* \end{bmatrix} = 0 \tag{35}$$

and note that the solution map $z^* := z(p)$ must be such that $\frac{\mathrm{d}r(z(p), p)}{\mathrm{d}p} = 0$. Evaluating the total derivative of $r(z(p), p)$ w.r.t. p, one obtains

$$\frac{\mathrm{d}z}{\mathrm{d}p} = - \left(\frac{\partial r}{\partial z} \right)^{-1} \frac{\partial r}{\partial p}, \tag{36}$$

Data	$\mathcal{P}(T_i)$	V_i	$\nabla_{T_i} V_i$	$\nabla^2_{T_i} V_i$	$h_i(T_i)$	$\nabla_{T_i} h_i(T_i)$	$\nabla_{T_i} h_i(T_i)$
# Floats	2	1	2	3	2^2	2	2

TABLE 1 The information which needs to be sent from a vehicle to the central node each iterate. The central node needs to send 2 floats to each vehicle: the current primal solution candidate $T_i^{(c)}$.

and, using the chain rule that

$$\frac{\mathrm{d}^2 z}{\mathrm{d} p^2} = -\left(\frac{\partial r}{\partial z}\right)^{-1} \left(\frac{\mathrm{d}}{\mathrm{d} p}\left(\frac{\partial r}{\partial p}\right) + \frac{\mathrm{d}}{\mathrm{d} p}\left(\frac{\partial r}{\partial z}\right)\frac{\mathrm{d} z}{\mathrm{d} p}\right). \tag{37}$$

Both (36) and (37) exist if $(\frac{\partial r}{\partial z})^{-1}$ exists, and the latter is guaranteed if Linear Independence Constraint Qualification (LICQ) and the Second Order Sufficient Conditions (SOSC) hold at the primal-dual solution of (33). The sensitivities required to evaluate (29), (30) and (32) can thereby be obtained by solving (36) for the KKT conditions of the QP (11) using the parameters T_i, and by solving (36) and (37) for the KKT conditions of the LPs (17), (18), using the parameter t_i^{in}.

Note that if the program (33) is solved using a second-order method, the solver performs iterations similar to

$$z^{(c+1)} = z^{(c)} - \alpha^{(c)} \left(\frac{\partial r}{\partial z}\right)^{-1} r. \tag{38}$$

This means that the solver will have factorized the matrix $\frac{\partial r}{\partial z}$ at the solution z, whereby the evaluation of (36) can be done at little additional computational cost. Consequently, if a second-order method is used to evaluate $V_i(T_i)$ and $h_i(T_i)$ in the SQP, the derivatives are cheap to compute.

3.3 | Schematic algorithm

A convergence for SQP applied to (7) is given in[25]. The SQP procedure is performed as follows

1. Initialize the primal-dual variables $v^{(0)} = (T^{(0)}, \mu^{(0)})$.

2. $\forall i$: Solve QP (11) and LPs (18), (17) to get cost, constraints and derivatives.

3. Assemble and solve QP to get $\Delta v^{(c)}$

4. Perform Projection Line search according to Section 3.1.2:

   (a) $\forall i$: Solve LPs (18), (17) at $T_i^\alpha = T_i^{(c)} + \alpha^{(c)}\Delta T_i^{(c)}$ and compute $\mathcal{P}_i(T_i^\alpha)$. Solve QP (11) with $\mathcal{P}_i(T_i^\alpha)$ to get V_i, ∇V_i, $\nabla^2 V_i$, h_i, ∇h_i, $\nabla^2 h_i$.

   (b) Assemble $M^\alpha = \sum_{i=1}^{N_a} V_i(\mathcal{P}_i(T_i^\alpha)) + \rho^{(c)}||h_s(\mathcal{P}(T^\alpha))\lambda\mathbb{Z}||_1$

   (c) If the Armijo condition (26) is not fulfilled, set $\alpha^{(c)} = \alpha^{(c)}\beta$, $\beta \in (0, 1)$ and repeat from 4a.

5. Perform primal-dual update (28a). If a solution is not reached, increment c and repeat from 3.

The computations in steps 2 and 4a consist of the solution of QP (11) and LPs (17),(18) and the associated derivative computations detailed in Section 3.2. We emphasize that most computations are separable and can be performed in parallel on board the vehicles. However, while the LPs (18)-(18) can be solved in parallel they must be solved before the the QP (11) due to the use of $\mathcal{P}_i(T_i)$. The non-parallelizable part of the algorithm is the formation and solution of the QP-subproblem (23), which thereby necessitates a central network node. In the scenarios considered in this paper, the Coordinator shown in Figure 1 a takes this role. Note that the SQP sub problem (23) has $2N$ variables and $5N - 1$ constraints, and will be significantly smaller than the vehicle-level QP (11) in moderately sized scenarios and horizons N_i. In such cases, the computational bottleneck will be the evaluation of $V_i(T_i^{(c)})$ and $h_i(t_i^{\mathrm{in}(c)})$ and the associated derivatives.

Communication Aspects

The information that needs to be communicated from the central node to the vehicles is only the currently held primal solution candidate, consisting of $2N_a$ floats. The data a vehicle is required to communicate to the central node is listed in Table 1 .

Consequently, each iterate will involve the communication of at least 14 floats per vehicle, and additional 14 for each reduction of the step size α. This will increase the time per iterate, and can, depending on communication protocol and implementation, constitute the bulk of the time required to solve the problem.

We note that it is only necessary to resend $\mathcal{P}(T_i^{\alpha})$ and $V_i(\mathcal{P}(T_i^{\alpha}))$ to evaluate the merit function M^{α} and check the Armijo condition (26). The remaining information in Table 1 could then be sent after the primal-dual update. However, while the amount of data transmitted would be less, two additional rounds of communication would be needed: one from the central node, notifying the vehicles of step acceptance, and one from the vehicles containing the remaining information. As practical communication systems include a significant overhead, the total time required by communication would be notably higher. While this requires the vehicles to compute the derivatives of $V_i(T_i)$ and $h_i(T_i)$ when not strictly needed, the calculations can be made highly efficient and will have a small impact on the total solve time (c.f. Section 3.2).

Besides what is mentioned here, the algorithm will require the communication of a number of logical variables, e.g. commands for algorithm start and stop, step acceptance etc. However, the amount of such data per iterate amounts to a few bytes and since it can be sent together with the other data, the additional time required is negligible.

3.4 | Real time implementation of the intersection-level controller

While the vehicle-level control loops consists of the on-line solution of QP (11), the on-line solution of NLP (7) forms the basis of the intersection-level control loop. In particular, when NLP (7) is solved using the SQP introduced in this section, it includes distributed computation and wireless communication. In the ideal case, algorithmic overhead, computation and communication require a negligible amount of time compared to the time scale of the system dynamics, and the current vehicle state $x_{i,k}$ (and therefore $V_i(x_{i,k}, T_i)$ and $h_i(x_{i,k}, T_i)$ can be considered constant during the time it takes to solve the SQP. However, non negligible delays can be expected in a real application, and the vehicle state might change significantly between the SQP iterations. The use of wireless communication in particular is a likely source of comparatively large delays, as packet drops are likely to occur and subsequent re-transmissions of data often are necessary.

This raises what is known as the real-time dilemma[30]: Should the NLP be solved to convergence when the resulting T will be outdated w.r.t. the system state, or should an approximate solution T be sought using the most up-to-date information. This implies that the resulting control law $T^*(x_{i,k})$ will be a sub-optimal approximation of the truly optimal solution feedback, regardless of how the dilemma is handled. In this paper, we consider two different solutions to the dilemma:

3.4.1 | Alternative 1: Solving the intersection-level NLP to convergence from a predicted state

In the first solution, which we denote the *Converged* controller, we use a scheme similar to that of[31] or[32], where $V_i(T_i)$, $h_i(T_i)$ are computed from a predicted future state $\bar{x}_{i,K}$ rather than the current state $x_{i,k}$, and the resulting control law $T^*(\bar{X}_{i,k})$ is not applied until time t_K. The predicted state is obtained from the open loop predictions of QP (11) in the vehicle-level control loops, which are computed using a previous timeslot schedule T. A block diagram of the scheme is illustrated in Figure 2 a.

If the evolution of the real system stays close to the predicted trajectories so that $\bar{x}_{i,K}$ predicted at some $k < K$ is close to the actual state $x_{i,K}$ at K, the resulting intersection-level control-law will $T^*(\bar{X}_{i,k})$ provide a good approximation to $T^*(X_{i,k})$. However, the scheme will introduce a significant delay in the feedback of the intersection-level scheme. In fact, denoting update frequency of the intersection-level controller t_s^T, we note that the reaction to a perturbation that occurs between kt_s^T and $(k+1)t_s^T$ will not be applied to the system until $(k+2)t_s^T$.

3.4.2 | Alternative 2: Approximate the intersection-level NLP in a RTI-fashion

The second solution, which we call the *1-step* controller, does not solve the SQP to convergence, and thereby avoids long solution times. Instead, we adopt a strategy where the current state measurements are used to compute $V_i(T_i)$, $h_i(T_i)$ and their derivatives, but where only one full SQP step is taken in the solution of (7). The resulting control law $T^{(1)}(X_k) = T^{(0)} + \Delta T^{(0)}(X_k)$ is thereafter applied directly to the vehicles. While the control law is approximate, it enables rapid feedback and reaction to perturbations to the vehicles. The procedure is illustrated in Figure 2 b.

This can be seen as an application of the Real Team Iteration (RTI) scheme for dynamic optimization[27], applied to the NLP (7). However, we note that the dynamic optimization problems (11) that gives the vehicle control commands are solved to convergence at all times.
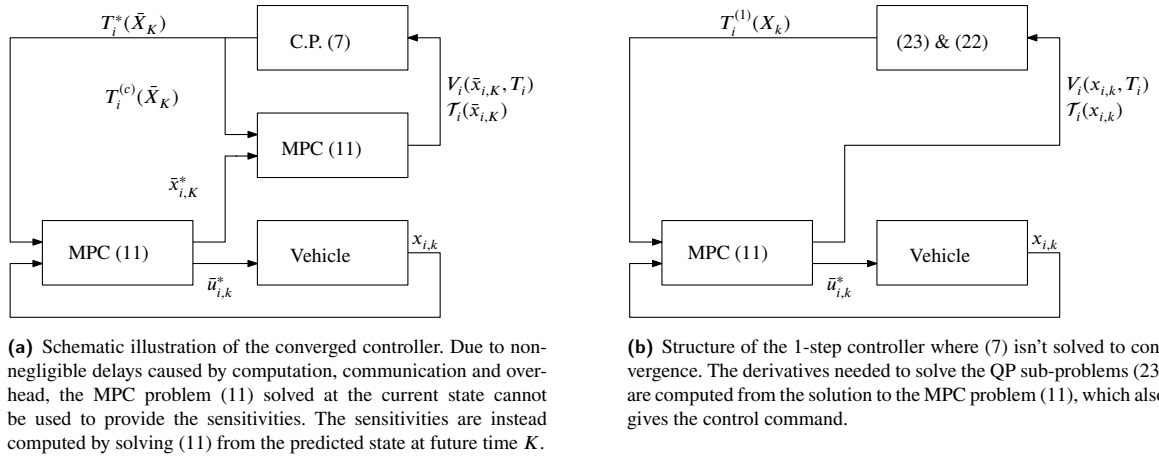
(a) Schematic illustration of the converged controller. Due to non-negligible delays caused by computation, communication and overhead, the MPC problem (11) solved at the current state cannot be used to provide the sensitivities. The sensitivities are instead computed by solving (11) from the predicted state at future time K.

(b) Structure of the 1-step controller where (7) isn't solved to convergence. The derivatives needed to solve the QP sub-problems (23) are computed from the solution to the MPC problem (11), which also gives the control command.

FIGURE 2 Schematic illustration of the two intersection-level controllers

4 | EXPERIMENTAL VALIDATION

In this section we describe an experimental setup which was used to validate the bi-level controller described in Section 2.3. In particular, we detail an implementation of the distributed SQP described in Section 3.1 in which most computations are performed on board the vehicles and communicated to a central coordinating unit using V2V communication. We also provide details on the hardware platform used.

4.1 | Practical implementation of the Vehicle-level Control-loop

The vehicle-level control loop consists of solving the problem (11) every time instant at the current state $x_{i,k}$, using the time-slot T_i and applying the resulting optimal $u_{i,k}$ to the vehicle. However, problem (11) differs from standard MPC formulations in that the position constraints (11e) and (11f) force the vehicle to be at a specific position at a given time. As the vehicle gets closer to the intersection, the ability to affect when the intersection is entered and departed diminishes. Moreover, in a real scenario, the closed-loop system is constantly exposed to perturbations in the form of plant-prediction model mismatches, measurement noise and other external disturbances. It therefore becomes increasingly likely that problem (11) is infeasible for $(x_{i,k}, T_i)$ when the vehicle gets closer to the intersection.

To address this issue we first relax the equality constraints (11e) and (11f) to the inequalities

$$p_i^{\text{in}} - p_i^d(t_i^{\text{in}}, w_i) \geq 0, \quad p_i^d(t_i^{\text{out}}, w_i) - p_i^{\text{out}} \geq 0. \tag{39}$$

With this relaxation, the vehicle is allowed to occupy the intersection within T_i rather than using the intersection throughout all of T_i. While this ensures that the controller, for instance, does not slow down the vehicle to stay longer in the intersection in response to a perturbation, infeasibilities are still possible. We therefore introduce an additional softening of the constraints (39) as

$$p_i^{\text{in}} - p_i^d(t_i^{\text{in}}, w_i) + \sigma_i^{\text{in}} \geq 0, \quad p_i^d(t_i^{\text{out}}, w_i) - p_i^{\text{out}} + \sigma_i^{\text{out}} \geq 0, \quad \sigma_i^{\text{in}} \geq 0 \quad \text{and} \quad \sigma_i^{\text{out}} \geq 0. \tag{40}$$

and add the term $P_i(\sigma_i) = \frac{1}{2}\sigma_i^\top \sigma_i \phi_i^q + \phi_i \mathbf{1}^\top \sigma_i$ to the objective, where $\sigma_i = (\sigma_i^{\text{in}}, \sigma_i^{\text{out}})$ and $\phi_i^q > 0$, $\phi_i > 0$ are penalty weights.

The relaxed-and-softened vehicle MPC problem solved at time t_k is thereby

$$\min_{\bar{w}_i, \sigma_i} \quad J_i(\bar{w}_i) + P_i(\sigma_i) \tag{41a}$$

$$\text{s.t.} \quad \bar{x}_{i,k} = x_{i,k} \tag{41b}$$

$$\bar{x}_{i,k+n+1} = A_i \bar{x}_{i,k+n} + B_i \bar{u}_{i,k+n}, \qquad n \in \mathbb{I}_{[0,N-1]}, \tag{41c}$$

$$D_i \bar{x}_{i,k+n} + E_i \bar{u}_{i,k+n}(t) \geq b_i,, \qquad n \in \mathbb{I}_{[0,N-1]}, \tag{41d}$$

$$p_i^{\text{in}} - p_i^d(t_i^{\text{in}}, \bar{w}_i) + \sigma_i^{\text{in}} \geq 0, \tag{41e}$$

$$p_i^d(t_i^{\text{out}}, \bar{w}_i) - p_i^{\text{out}} + \sigma_i^{\text{out}} \geq 0, \tag{41f}$$

$$\sigma_i^{\text{in}} \geq 0, \quad \sigma_i^{\text{out}} \geq 0. \tag{41g}$$

Here, we differentiate the state and control of the vehicle $x_{i,k}, u_{i,k}$ from the open loop predictions $\bar{x}_{i,k}, \bar{u}_{i,k}$. The control command applied at time t_k is the optimal open loop control command $u_{i,k} = \bar{u}_{i,k}^*$.

The softening of the constraints ensures that there will be no feasibility issues due to the position constraints (39). In fact, T_i no longer affects the feasibility of the optimization problem (41). Note that if ϕ_i is chosen large enough, $P_i(\sigma_i)^3$ is a so-called exact penalty function[28]. A well known property of exact penalty functions is that the problem with softened constraints will return a solution with $\sigma_i = 0$ whenever such a solution exists (see e.g.,[33] Theorem 14.3.1). Moreover, when ϕ_i is chosen large enough and no solution exists for $\sigma_i = 0$, the solution minimizes $||\sigma_i||_\infty$ and thereby the violations of the constraint (39)[23].

4.2 | Efficient Solution of the Quadratic and Linear Programs

General purpose solvers are too slow to give real-time feasible solutions to the vehicle-level QP (41) and the LPs (18), (17), and solvers tailored to the special structure of these optimization problems have to be considered.

As discussed in Section 3.2, in case second-order optimization methods are used to solve these QPs/LPs, the sensitivities of the cost function and constraints can be easily and cheaply computed from the local optimization problems by reusing the KKT matrix factorization available from the QPs/LPs solver. In this work, we used a version of the interior-point method (IPM) HPMPC[26] tailored to allow the efficient computation of the tangential predictor at the solution.

HPMPC provides an implementation of a Mehrotra's predictor-corrector IPM tailored for the solution of QPs in the form of optimal control problems (OCP). The IPM employs a backward Riccati recursion for the efficient computation of the search direction. As its linear algebra framework, HPMPC makes use of BLASFEO[?], which provides a set of linear algebra routines tailored to provide high computational performance for rather small matrix sizes, as typical in embedded optimization applications.

On the algorithmic side, the IPM in HPMPC is coupled with a partial condensing algorithm. Partial condensing[34] is a technique that allows to control the level of sparsity of an OCP problem by trading off horizon length with input vector size, by condensing block-wise the original OCP. It is possible to compute the theoretical optimal horizon length based on the analysis of the flop count of the algorithm. In practice, however, other factors affect the optimal choice of the horizon length, such as the performance of linear algebra routines[?]. The QP (41) is a perfect example of that. Since the state and input vector sizes are very small and the horizon length is long, partial condensing gives a QP reformulation that HPMPC can solve much faster, since many operations on small matrices (where the linear algebra performs poorly) are replaced with few operations on large matrices (where the linear algebra gives higher computational performance).

In this work, HPMPC has been modified to allow the efficient computation of sensitivities. Namely, the solver now allows the reuse of the last KKT matrix factorization (where Lagrange multipliers and slack variables of inequality constraints are frozen at their value close to the solution) to cheaply compute the solution of other systems of linear equations with different right hand side. If there are no changes in the active set, this allows the efficient computation of the tangential predictor around the current solution[35]. Therefore, the sensitivities in 36 can be cheaply computed by performing the partial condensing of the right hand side and the solution of the KKT system reusing the cached KKT matrix factorization. The computational cost of these operations is negligible with respect to the QP/LP solution, which comprises a complete partial condensing pre-processing step, plus a KKT matrix factorization and two KKT system solutions per IPM iteration (which are typically in the range of 6-15 per QP/LP solution).

---

[3]The quadratic term in $P_i(\sigma_i)$ is added for numerical reasons, and the parameter ϕ^q is typically small.

Efficient solution of the intersection-level problems

As noted in Section 3.3, the computational bottleneck of the SQP in moderately sized scenarios is the solution of the vehicle-level QP (41) and the LPs (18), (17). This is due to the comparatively small size of the QP subproblem (23) in such cases. General purpose QP solvers can therefore be fast enough and used to solve (23) in real-time. Due to this, MATLABs QP-solver quadprog was used in the experimental validation.

However, with an increasing number of vehicles, the time required to solve the QP subproblem (23) with a general purpose solver will approach that required by HPMPC for the solution of (41) and the LPs (18), (17). For large scenarios, solving (23) could therefore become a computational bottleneck of the SQP. It is therefore desirable to use efficient, structure exploiting solvers also for the QP subproblem (23). For this reason, we propose the following reformulation of (23):

$$\min_{\Delta \boldsymbol{T}, u^t} \quad \sum_{i=1}^{N_a} \frac{1}{2} \Delta T_i^\top H^{(c)} \Delta T_i + \nabla_{T_i} f(T_i^{(c)})^\top \Delta T_i \tag{42a}$$

$$\text{s.t.} \quad t_i^{\text{in}(c)} + \Delta t_{i+1}^{\text{in}} = u_i^t \tag{42b}$$

$$u_i^t - t_i^{\text{out}(c)} - \Delta t_i^{\text{out}} \geq 0 \tag{42c}$$

$$h_i(T_i^{(c)}) + \nabla_{T_i} h(T_i^{(c)})^\top \Delta T_i \geq 0. \tag{42d}$$

That is, the time-slot increments ΔT_i are formulated as states in the dynamical system (42b), where the variable $u^t = (u_1^t, \ldots, u_{N_a}^t)$ is introduced as a fictitious control. The precedence constraint (7c) is formulated as the path constraint (42c). The problem is thereby written on a stage-wise form for which efficient solvers as HPMPC can be deployed, and significant performance gains can be made. For instance, the typical time required to solve (23) for a three vehicle scenario with quadprog is around 2 ms using a standard laptop. The time required by HPMPC for the same problem using the same hardware lies around 40 μs. Moreover, the time-complexity of HPMPC is linear in the number of stages, and approximately 100 μs are required for a 30-vehicle scenario, where quadprog requires 7 ms. For a 300-vehicle scenario HPMPC requires approximately 850 μs to converge, while quadprog requires approximately 1 s.

4.3 | Algorithm

The procedures executed by the coordinator and the vehicles when the converged intersection controller is used are summarized in Algorithm 1 and Algorithm 2. Before the SQP is solved the first time, the central node requests the vehicles' non-coordinated optimal time-slot T_i which serves as the primal initial guess. In subsequent solutions of the SQP, the previous optimal solution is used instead. The solution is considered found when either $||(\nabla \mathcal{L}(z), h^-(\boldsymbol{T}))||_\infty < \varepsilon$ or when $||(\Delta \boldsymbol{T}, h^-(\boldsymbol{T}))||_\infty < \varepsilon$. On the vehicle side, the MPC problem (41), is solved every t_s using the most recently commanded time-slot T_i^*. When a request from the central node is registered, the vehicles evaluate and send the functions and derivatives needed to solve the SQP.

Implementation Restrictions and Practical Considerations

Due to implementation related details, the algorithm was executed with the period $t_s = 0.1$ s. This restricted the communication of new information from both the central node and the vehicles to occur at a maximum of 10 Hz. With an ideal communication system, one iteration of the SQP (with $\alpha^{(c)} = 1$) therefore requires t_s s to broadcast $\boldsymbol{T}$ (line 11 of Algorithm 1) and t_s s for the vehicles to respond with $\mathcal{D}_i(\mathcal{P}_i(T_i^\alpha))$ (line 12 of Algorithm 1), i.e. the lowest time required per SQP iterate is $2t_s$ s. Moreover, for the Converged intersection controller, each execution of the SQP commences t_s^T s before the resulting time-slots should be applied to the vehicles, which gives the algorithm t_s^T s to converge and notify the vehicles of the results. For simplicity, the period of the intersection-level control loop is set to t_s^T s, which due to the long expected solve times is set to $t_s^T = 3$ s. Finally, to ensure that all vehicles are predicted to be before the intersection when the new time-slots are applied, the SQP is only solved when no vehicle is close to the intersection. In particular, for scenarios where the desired speed is $v^{\text{ref}} = 50$ km/h, the SQP is suspended when the first vehicle is 50 m away from the intersection. We emphasize that this modification is done for simplicity of implementation and that the problem formulation allows the SQP to be solved with a vehicle to be inside the intersection.

Algorithm 1 Central Node. Here, $\mathcal{D}_i(T_i) = \{V_i(T_i), \nabla_{T_i} V_i(T_i), \nabla^2_{T_i} V_i(T_i), h_i(T_i), \nabla_{T_i} h_i(T_i), \nabla^2_{T_i} h_i(T_i)\}$, where the dependence on the initial state $\hat{x}_{i,0}$ has been dropped for notational simplicity.

1: Send coordination start state x_i^{start} and start time t^{start}.
2: Request $T_i^{(0)}$ from vehicles, initialize $\mu^{(0)}$, $K = 0$
3: **loop**
4:      Wait until $t == t^{start} + (K-1)t_s^T$
5:      Set $t_K = K t_s^T$, $T = T^*$, $\mu = \mu^*$, $r = 0$
6:      Broadcast T, t_K, r and request $\mathcal{D}_i(T_i)$ computed at t_K.
7:      Wait until all vehicles has responded.
8:      **while** exit conditions not fulfilled **do**
9:          Set $r = r + 1$, Assemble and solve (23) for Δz, set $\alpha^{(c)} = 1$
10:         **loop**
11:             Broadcast $T_i^\alpha = T_i^{(c)} + \alpha^{(c)} \Delta T_i^{(c)}$ and request $\mathcal{D}_i(T_i^\alpha)$ and $\mathcal{P}_i(T_i^\alpha)$
12:             Wait until all vehicles have responded.
13:             $M^\alpha = V(\mathcal{P}(T^\alpha)) + \rho^{(c)} ||h_s(\mathcal{P}(T^\alpha))\lambda \mathbb{Z}||_1$
14:             **if** $M^\alpha \leq M(T^{(c)}) + \gamma D_{\Delta T^{(c)}} M(T^{(c)})\alpha^{(c)}$ **then**
15:                 $T^{(c+1)} = \mathcal{P}(T^\alpha)$
16:                 $\mu^{(c+1)} = \mu^{(c)} + \alpha^{(c)}(\tilde{\mu}^{(c)} - \mu^{(c)})$
17:                 Exit loop.
18:             **else**
19:                 $\alpha^{(c)} = \alpha^{(c)}\beta$
20:             **end if**
21:         **end loop**
22:     **end while**
23:     Set $T^* = T^{(c)}$ and send out time-slot to apply T^*.
24:     $K \leftarrow K + 1$
25: **end loop**

Algorithm 2 Vehicle

1: **loop**
2:      Estimate current state $x_{i,0}$, Get synchronized time t
3:      **if** Central node sends new time-slot to apply **then**
4:          Receive T_i^*
5:          $T_i^{local} = T_i^* - t$
6:      **end if**
7:      Solve (41) with $(x_{i,0}, T_i^{local})$
8:      Apply optimal $u_{i,0}$ to vehicle
9:      **if** Central node request $\mathcal{D}_i(T_i)$ **then**
10:         Receive T_i, t_K, r
11:         $\tilde{T}_i^{local} = T_i - t_K$
12:         **if** r == 0 **then**
13:             Store $x_{i,K}$ from prediction at $t_K - t$ computed on Line 7
14:             Solve LPs for $L^{in}(x_{i,K})$ and $U^{in}(x_{i,K})$
15:         **end if**
16:         Solve LPs (18), (17) and evaluate $h_i(\tilde{T}_i^{local})$, $\nabla_{T_i} h_i(\tilde{T}_i^{local})$, $\nabla^2_{T_i} h_i(\tilde{T}_i^{local})$.
17:         Compute $\mathcal{P}_i(\tilde{T}_i^{local})$ through (27).
18:         Solve QP (11) with $(x_{i,K}, \mathcal{P}_i(\tilde{T}_i^{local}))$ and compute $V(x_{i,K}, \mathcal{P}_i(\tilde{T}_i^{local}))$, $\nabla_{T_i} V(x_{i,K}, \mathcal{P}_i(\tilde{T}_i^{local}))$ and $\nabla^2_{T_i} V(x_{i,K}, \mathcal{P}_i(\tilde{T}_i^{local}))$.
19:         Send $\mathcal{P}_i(\tilde{T}_i^{local})$ and $\mathcal{D}_i(\mathcal{P}_i(\tilde{T}_i^{local}))$ to central node.
20:     **end if**
21: **end loop**

4.4 | Test Platform

The coordination controller was tested at the Asta Zero proving ground outside Gothenburg, Sweden. The test platform consisted of the three different Volvo vehicles shown in Figure 3 a: One Volvo S60 T5 Petrol turbo sedan, one Volvo S60 D5 Turbo Diesel sedan and one Volvo XC90 T6 Petrol turbo SUV. All cars were equipped with automatic gearboxes and an interface for external control of the longitudinal motion. In particular, all vehicles were commanded by supplying a desired longitudinal acceleration to a controller, which thereafter sent the appropriate commands to the engine, gear-box and friction brakes. The vehicles had an on board senor suite consisting of wheel encoders, inertial measurement units and Real-Time Kinematic (RTK) GPS receivers. The latter was capable of providing positioning estimates where the measurement error standard deviation were as low as $\sigma_{GPS} = 0.05$ m as well as global-time synchronization. To improve the positioning estimates and handle sporadic GPS outages, Extend Kalman Filters based on that presented in[36] were used in all vehicles to fuse the available sensor data. The one-dimensional position $p_{i,k}$ was constructed by first projecting the estimate of the global-position onto a reference path along the road and then taking $p_{i,k}$ as the geodesic distance along the path from the projected point to the start of the intersection.

Moreover, each vehicle was equipped with ITS G5 compliant vehicle-to-vehicle communication equipment from RENDITS[37]. On each vehicle, the experiment software ran on two computational units: one MicroAutoBox II (MABx) real-time prototyping platform, interfacing with the vehicle, sensors and communication equipment which ran the algorithm logic and state estimation, and one laptop on which the QPs and LPs were solved. The hardware setup in one of the vehicles is shown in Figure 3 b.

FIGURE 3 Photos of the experimental hardware platform. (a) Shows the three vehicles used, (b) Shows the hardware installation in one of the vehicles. The 3G router was used to provide an IP link for the RTK corrections used by the GPS receiver

Finally, the experimental setup included a central coordinating unit which consisted of a laptop and a vehicle-to-vehicle communication device. The central coordinating unit was used to control and monitor the experiments and executed the central parts of the SQP.

Prediction model, Objective and Parameters

The prediction model used during the experiments was a simple double integrator, $\dot{p}_i(t) = v_i(t)$, $\dot{v}_i(t) = u_i(t)$, where the acceleration is the input and $x_i(t) = (p_i(t), v_i(t))$. For this dynamical system $p_i^d(t, w_i) = p_{i,k} + (t - kt_s)v_{i,k} + \frac{1}{2}(t - kt_s)^2 u_{i,k}$, with $k = \text{floor}(t/T_s)$. Furthermore, we employed the objective function

$$J_i^d(w_i) = (v_{i,N} - v_i^{\text{ref}})^2 Q_i^f + \sum_{i=0}^{N-1} (v_{i,k} - v_i^{\text{ref}})^2 Q_i + R_i u_{i,k}^2, \tag{43}$$

where the desired speed v_i^{ref}, and the objective function weights $Q_i > 0$ and $R_i > 0$ were varied between different instances of the experiment. The state and control were constrained to $v_{i,k} \geq 0$ and $u_{i,k} \in [-4, 1.6]$ m/s^2, where the latter was due to limitations in the vehicle actuation interfaces. The vehicle-level control loops were closed with $t_s = 0.1$ s and the horizon length was set to $N = 200$. The objective and prediction model were chosen due to their simplicity. In particular, the dynamics does not include any parameters to identify, and the objective enables an intuitive understanding of how the solution will change with variations to the penalty weights. However, we want to emphasize that these choices are not restrictive and that other linear-quadratic models are possible. For instance, in[?] a prediction model based on a linearization of a non-linear vehicle model is used together with a quadratic approximation of an economic objective function obtained through the method presented in[38].

5 | RESULTS

In this section, we present and discuss results from the experimental campaign, which demonstrates the performance of the semi-distributed SQP and both the converged and 1-step intersection-level controllers. In total, more than 80 experiments were performed, where the initial conditions, objective function weights and other parameters were varied. In all experiment instances, the vehicles were first controlled to a predefined starting state, typically one where a collision would occur if no action was taken, before the bi-level controller was initialized. The experiments were performed in two different modes: in an actual intersection, as illustrated in Figure 4 a and in a parallel configuration where the approaching roads were laid out next to each other and the intersection was represented by a segment on the road, as illustrated in Figure 4 b. The latter was used to enable evaluation of the controller without a risk of collision, and is the primary source of the data reported in this section. However, the interested reader can find video material from experiments performed in the crossing configuration at[39].

For comparison, some simulation results of the closed-loop system are also provided in this section. In these cases, the vehicles were simulated with the nominal model.

(a) Aerial photo of the crossing configuration used in the experimental validation. The white lines before the square representing the intersection illustrates the different safety margins employed.

(b) Photo of the parallel configuration used in the experimental validation. The white lines mark the beginning and end of the intersection, p_i^{in} and p_i^{out}, and collisions are thereby avoided when only one vehicle is between the two white lines at a given time

FIGURE 4 Photos of the two different configurations used in the experimental validation.

5.1 | **Evaluation of the semi-distributed SQP**

In this subsection we present and analyze data from the implementation of the semi-distributed SQP used in the experiment. We focus on a experimental scenario with the objective function weights $Q_i = 1$, $R_i = 10$, the reference speed $v_i^{\text{ref}} = 50$ km/h, $\forall i$, and where all vehicles are initiated at $p_{i,0} = -200$ with $v_{i,0} = 50$ km/h. Data from the first SQP instance, where the problem is solved from $x_{i,0} = (p_{i,0}, v_{i,0})$, $\forall i$, is presented in Figure 5 . In particular, Figure 5 a and Figure 5 b show the solutions to the vehicle-level problems corresponding to the iterates $T^{(c)}$ of the SQP. The two horizontal lines in Figure 5 a represent the beginning and end of the intersection, i.e. collisions are avoided when at most one trajectory is between the lines at all times, and the primal iterates $T^{(c)}$ are shown as vertical lines.

In this case, the solver is initialized at the uncoordinated solution, where all vehicles keeps v_i^{ref} and therefore occupy the intersection simultaneously. As can be seen in Figure 5 a, the time-slots T satisfy the order constraints $h_i(T) \geq 0$ already after the first iterate, whereby a collision free solution is available. The subsequent two iterations retains feasibility and improve the solution. Full steps ($\alpha = 1$) are taken in all iterates.

The algorithm progress measure shown in Figure 5 c further illustrates this fact: the feasibility is reached after the first iterate. Note the relatively loose convergence threshold $\varepsilon = 10^{-3}$, which is selected in relation to the properties of the physical system. In particular, the GPS provides measurements with a positioning error standard deviation around $\sigma_{\text{GPS}} = 0.05$ m. The standard deviation of the error between the commanded time-slot T_i and that resulting from the closed-loop application of (11) using T_i will therefore be above one millisecond[4] for speeds around 50 km/h. Enforcing constraint satisfaction or changes to the primal variables below 10^{-3} $[s]$ will consequently have no noticeable effect on the physical system.

The small impact of small changes in T_i on the vehicle trajectories is further illustrated in Figure 5 a and Figure 5 b, where the solutions of (11) corresponding to iteration 2 and 3 are indistinguishable. In fact, the difference in $\bar{u}_{i,0}$ between iteration 2 and 3 is below the resolution of the actuation interface, i.e. the applied control command $\bar{u}_{i,0}$ after iteration 2 and 3 would be interpreted as identical by the vehicle.

Solution time analysis

The mean computation time was 1.664 ms with $\sigma \approx 1.263$ ms ($n = 52736$) for the LPs, and 1.607 ms with $\sigma \approx 1.107$ ms ($n = 246132$) for the vehicle-level QPs . Since the SQP normally only required a few iterations to converge to a relevant threshold, a solution to (7) should have been found within a few hundreds of a second. However, due to a rudimentary implementation and hardware limitations this was not the case in during the experiments. Instead, the average time required to solve the SQP was 1.740 s with $\sigma \approx 0.406$ s ($n = 130$).

An example of how this time is spent is provided in Figure 6 which shows a timeline from the SQP instance shown in Figure 5 . In Figure 6 , the width of the bars represent the time spent solving the LPs (18), (17) and the QP (11) for the

---

[4]This is true for the otherwise ideal case. In realty, other uncertainties are present as well and the performance is in practice therefore even worse.

FIGURE 5 Example of the progression of the SQP algorithm in one instance from the experimental campaign. (a) Shows the primal iterates $T_i^{(c)}$ as vertical bars and the position trajectory $\bar{p}_{i,k}$ for the solutions to (11) corresponding to the primal iterates, (b) shows the corresponding velocity $\bar{v}_{i,k}$ and input $\bar{u}_{i,k}$ sequences and (c) the corresponding algorithm progress measure, where the dashed line is the termination tolerance. The data from the different vehicles in (a) and (b) are differentiated by color.

vehicles, corresponding to Lines 16 and 18 of Algorithm 2, and the time spent solving the QP-subproblem 23 in the central node, corresponding to Line 9 of Algorithm 1. As the figure illustrates, only 55 ms, corresponding to 1.6% of the total time 2.723 s, is spent on computation. The time required for the other operations relevant for the SQP in Algorithms 1 and 2 is negligible, and the remaining 98.4%, i.e. the white gaps of Figure 6 a, is primarily spent in the waiting states of Lines 7 and 12 in Algorithm 1. Possible explanations for the long waits are inefficient buffer handling in the communication modules of the MABx and packet drops in the wireless links. To a smaller extent, the delays are due to the low communication frequency used and the lack of synchronization between the vehicles and the central node, discussed in Section 4.3. The impact of the slow update rate can be seen in a comparison between the computations of vehicle 1 and 2: in iteration 1, occurring around $t = -1.5$ s in Figure 6 a, vehicle 1 and 2 perform their computations simultaneously. in iterations 0, 2 and 3, on the other hand, there is delay of $t_s = 0.1$ s between the vehicles. The explanation is that a message is processed by Algorithm 2 at a time $(k+1)t_s$ or kt_s depending on when it is received relative to the ticks of the local clock. Small variations in the the reception time can therefore cause a variation of $t_s = 0.1$ s in the relative time between the processing of a message in two cars. Moreover, the lack of synchronization can be observed when vehicle 2 and 3 are compared: the computations of vehicle 3 consistently occur around 0.07 s before those of vehicle 2. We want to highlight that the long time required to solve the SQP therefore is almost entirely related to our specific implementation, and that it is likely that performance would improve dramatically with a few improvements.

For comparison, we provide an example of the time-line for a more efficient implementation of the SQP in Figure 6 b. In this case, the long and unnecessary waits have been removed and the algorithm is synchronized between the vehicles, but everything else kept unchanged. The time marked as required by communication is taken from[36], where an empirical study is presented on the time performance of the communication system used during the experiment. With such an implementation, the same problem instance would be solved in 0.074 s, where 48% would be spent on communication, 38.5% on the solution of the vehicle-level LPs and QP, and 13.5% on solving the QP-subproblem (23).

Finally, it is reasonable to expect even lower solution times with improvements to both problem formulation and equipment. For instance, the use of the re-formulation of the QP-subproblem and its solution with HPMPC would, as discussed in 4.2, significantly reduce the time required for the central computations. Furthermore, solving the two LPs (18), (17) in parallel on each vehicle would shorten the time required for the vehicle-side computations. Finally, while the general purpose V2V equipment requires around 4 ms for each transmission, a tailored communication protocol could be made significantly faster. For instance, it is reported in[40] that the time to transmit data using the 802.11 p physical layer could be as low as $t_{\text{com}} = 40 + \text{ceil}\left((n_{\text{data bits}} + 22)/48\right) 8$ μs. Since all vehicles could use different channels and transmit their data to the central node simultaneously, sending 12 floats per vehicle in double precision would thereby take 176 μs. The time-line for a solution of the same SQP instance in this hypothetical setting is shown in Figure 6 c, where the solution would be found in 0.0256 s.
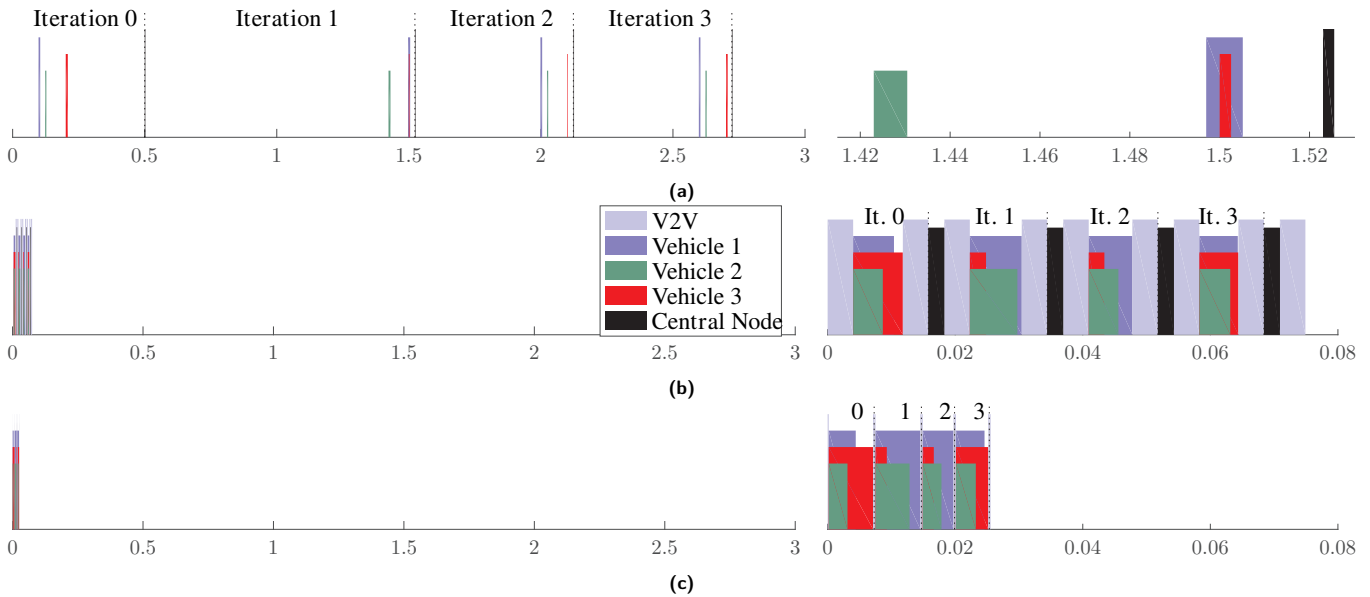
FIGURE 6 Illustration of the time spent on computation in the of the SQP instance shown in of Figure 5 , The width of the bars corresponds to the time spent for the different parts of Algorithms 1 and 2. For the vehicles, it consists of the time required to solve (18), (17) and (11) sequentially, and for the central node it consists of the time required to solve (23) with MATLABs quadprog. The bars are shown with different heights to ease visualization. Subfigure (a) shows the timings recorded during the experiment, (b) shows the likely timings of an improved implementation and (c) illustrates the timings of an hypothetical ideal implementation which uses a tailored communication protocol.

5.1.1 | Consistency

The algorithm consistently exhibited the same behavior as in Figure 5 , with almost immediate feasibility followed by a few optimality-improving iterations. The algorithm progress measures of a selection of the SQP instances are provided in Figure 7 for illustration. In particular, the convergence to ϵ was achieved in 2 iterations in 7% of the instances, in 3 iterations in 50%, in 4 iterations in 28% and in 5 iterations in 15%.

Finally, we have noticed that reduced steps ($\alpha < 0$) are required only in "hard" scenarios. Examples of "hard" scenarios include those where the coordination is initiated very close to the intersection for the initial vehicle velocity, and those where a large number of vehicles need to cross the intersection simultaneously. For practical reasons we were not able to perform sufficiently hard scenarios during the experimental campaign, due to which full steps were taken in all experimental SQP instances. For illustration purposes, the progress measures of a 3-vehicle scenario where reduced steps were taken is given in Figure 7 c. This scenario was particularly hard, since $p_{1,0} = -45$ m, $v_{i,0} = 45$ km/h while $p_{2,0} = p_{3,0} = -40$ m, $v_{2,0} = v_{3,0} = 50$ km/h, such that the first and last vehicles were forced perform very aggressive maneuvers to avoid collision.

5.1.2 | A Larger Example

To further demonstrate the behavior of the semi-distributed SQP, we present simulated results from a larger problem instance in Figure 8 . In this scenario, 12 vehicles are randomly generated at distances between 50 and 200 m from the intersection at 50 km/h. As Figure 8 b illustrates, the algorithm exhibits the same behavior as in the smaller scenarios: feasibility, and thereby collision avoidance is reached rapidly, in this case after the second iterate, and thereafter small adjustments towards optimality are performed.

While the implementation of the SQP used in the experiments would require prohibitively long time to solve the problem, it would be solved in 0.224 s with the improved implementation of Figure 6 b, and attain feasibility in 0.037 s. Moreover, the problem would be solved in 0.089 s with the ideal implementation discussed in Section 5.1, where a feasible solution would be available in 0.015 s.

FIGURE 7 Subfigure (a) and (b) shows the development of the progress measure in a selection of SQP instances from the experimental campaign. In both figures, the stars indicates cases when the algorithm was terminated due to only one progress measure reaching the threshold. Subfigure (c) shows the progress measure in a case where reduced steps are taken. Here, black plot show α, and the remaining colors follow Figure 5 c. In (a), (b) and (c), the dashed line is the tolerance $\varepsilon = 10^{-3}$.

FIGURE 8 Example scenario consisting of 12 cars in a four-way intersection. The positions of the vehicles are shown in (a), where the different colors differentiate different lanes. In (b), the algorithm performance measures are displayed together with the cumulative number of floating point numbers passed in the two-way communication between the central node and the vehicles.

5.2 | Evaluation and comparison of controllers

In this subsection we present and analyze both simulated and experimental data on the performance of the bi-level closed-loop controller In particular, we provide comparisons between the converged and 1-step formulations of the intersection-level controller and study their ability to reject large perturbations.

5.2.1 | Vehicle Level Control Loop

Regardless of how the time-slots T are computed, the accuracy with which the vehicles conform to the time-slots determines whether or not the closed-loop system is collision free. Even though a very simple dynamic model was used in the vehicle-level MPC during the experiment, the ability of the closed-loop system to satisfy the position constraints (41e) and (41f) was remarkable. In particular, even though several non-modeled nonlinearities were present the difference between the commanded (T_i) and actual (T_i^{Actual}) time-slot was small in most cases. A total of 450 experimental evaluations of this difference were performed, consisting of both $t_i^{\text{in Actual}} - t_i^{\text{in}}$ and $t_i^{\text{out Actual}} - t_i^{\text{out}}$ for all three vehicles and all experiment runs. The error was on the dangerous side, i.e. when the vehicle enters the intersection too early ($t_i^{\text{in Actual}} < t_i^{\text{in}}$) or when it leaves the intersection too late

FIGURE 9 Illustration of the vehicle-level MPC performance. Statistics on the difference between the commanded and actual t_i^{in} and t_i^{out} is shown in (a). In particular, only dangerous violations such that $t_i^{\text{in Actual}} < t_i^{\text{in}}$ or $t_i^{\text{out Actual}} > t_i^{\text{out}}$ is considered. The corresponding position trajectories are shown in (b) and (c). Trajectories that satisfy and violate (41e) and (41f) are colored green and red respectively. The gray area in (a) and (b) corresponds to the constraints (41e) and (41f), and are consequently not crossed by safe trajectories. The data is obtained from all three vehicles during 75 experiments and all trajectories are shifted in $p_{i,k}$ and t to ease visualization.

$(t_i^{\text{out Actual}} > t_i^{\text{out}})$, in 43.3% of the cases. The distribution of the error is given in Figure 9 a, and we emphasize that for more than 90% of the potentially dangerous constraint violations, the errors were below 0.03 s. To illustrate how this small error translates to collision risks, the corresponding trajectories in the time-position space is given in Figure 9 b and Figure 9 c for t_i^{in} and t_i^{out} respectively. As can be seen in Figure 9 c, the magnitude of the violation is such that a few decimeters of the leaving vehicle is inside the intersection when the next vehicle enters. A constraint tightening is discussed in [23] using which collision avoidance can be guaranteed for the closed-loop system even with potentially dangerous constraint violations.

As shown in Figure 9 c, most dangerous constraint violations occur for the out time constraint (41f), i.e, they are such that the vehicles leaves the intersection too late. This is likely due to the (non-modeled) actuator dynamics being faster in deceleration than in acceleration (c.f. friction brakes and internal combustion engine). Successful compensation of errors due to measurement noise and prediction model inaccuracies can therefore be made closer to the intersection when these cause the vehicle to enter early rather than leave late. It is expected that all constraint violations could be decreased by using a more sophisticated prediction model in (41), more accurate sensors and a higher update frequency t_{s}.

5.2.2 | Intersection Level Control Loop

To enable comparison, the same experiments were performed using both the converged and 1-step controller. Data from the application of both controllers to one experimental scenario is presented in Figure 10 , where the scenario parameters were as in Section 5.1. As the figure illustrates, the difference between the two controllers is very small: the acceleration and velocity profiles in Figure 10 b show a high degree of similarity, and the position trajectories in Figure 10 a are almost indistinguishable. The differences are perhaps most clearly seen in the lower plot of Figure 10 c, which shows the changes in t_i^{in} compared to the first coordinated solution at $t = 0$. Here, the smaller but more frequent changes to $\boldsymbol{T}$ by the 1-step controller are clearly differentiated from the less frequent but larger adjustments performed by the converged controller. Note that, while updates are more frequent for the 1-step controller, they are still significantly slower than the vehicle-level update frequency of 10 Hz. The reason is that each adjustment first requires all vehicles to send the relevant information to the central node, which thereafter can solve the QP subproblem (23) and send the updated time-slot back to the vehicles. The process thereby involves the same type of waiting and delays as discussed in Section 5.1.

Note also that for both controllers, the initial time-slot schedule $\boldsymbol{T}$ is continuously pushed to later times. This is likely caused by inaccuracies in the prediction model which cause the real system to lag slightly behind the predictions. This explanation is consistent with the nature of the constraint violations discussed in Section 5.2.1, in particular those shown in Figure 9 c. The use of a more accurate prediction model would likely affect the behavior.

We want to highlight that the magnitude of the input commands and acceleration as well as the resulting changes in velocity are all small. For comparison, experiments have shown that human drivers decelerate with down to -1.9 m/s^2 during intersection approaches without stops (light switching from red to green) and down to -4.5 m/s^2 for solid red lights [41].

FIGURE 10 Data from two experimental runs using the 1-step and converged intersection-level controllers, where $Q_i = 1$, $R_i = 10$, $v_i^{\text{ref}} = 50$ km/h, $\forall i$. The position trajectories are shown in (a) and the velocity and acceleration trajectories in (b). The upper plot of (c) shows the output of the vehicle-level MPC controllers whereas the lower plot illustrates the output of the intersection-level controller. To promote visibility, the lower plot in (c) shows only changes in t_i^{in} with respect to the initial schedule t_i^{in0}, which is the same for both controllers. For both controllers, the time-slot schedule is frozen at $t \approx 10$ s. Due to substantial noise levels, the acceleration data in (b) has been smoothened with a non-causal filter to promote visibility. Finally, the noise in the control signals of vehicle 2 and 3 seen in (c), is due to a GPS problem. The issue is thoroughly discussed in [23].
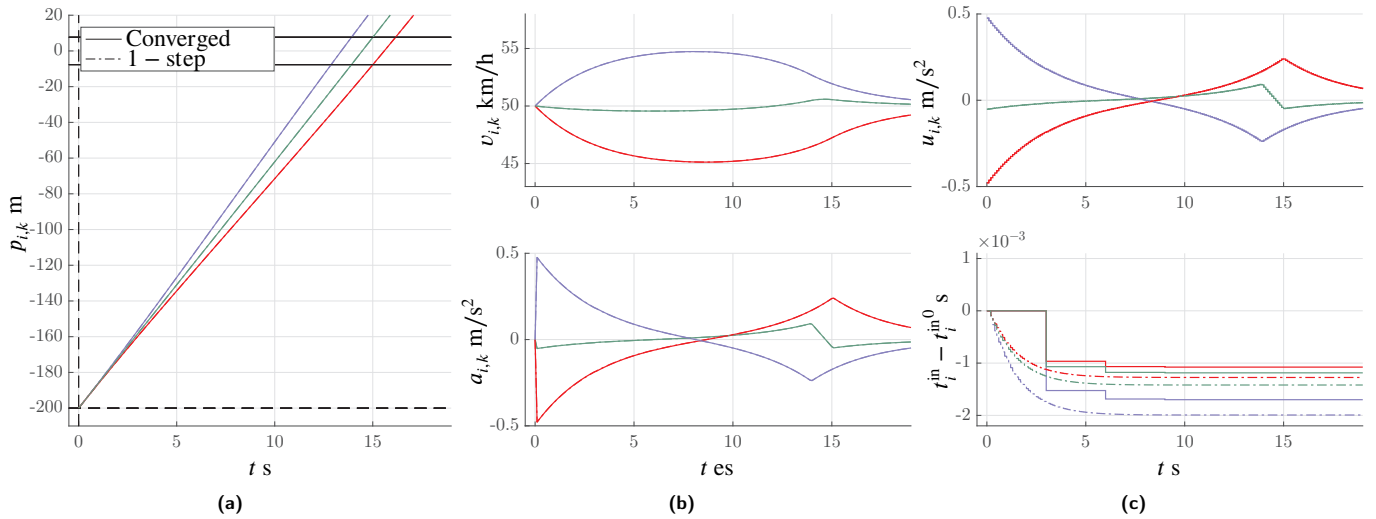
FIGURE 11 Results from simulations where the converged and 1-step intersection controllers were used. Plots as in Figure 10 .

To illustrate the effects of using a simple implementation and deploying the controller on a real system, results from a simulation of both controllers and the same scenario is provided in Figure 11 . Since one iteration of the SQP requires less than 20 ms in the ideal case, including 2-way communication, the 1-step updates are computed with the same frequency as the vehicle-level controllers (0.1 s) in simulation. The differences between the controllers are not surprisingly even smaller in the simulated case, and is again seen most clearly in the control output plot of Figure 11 c.
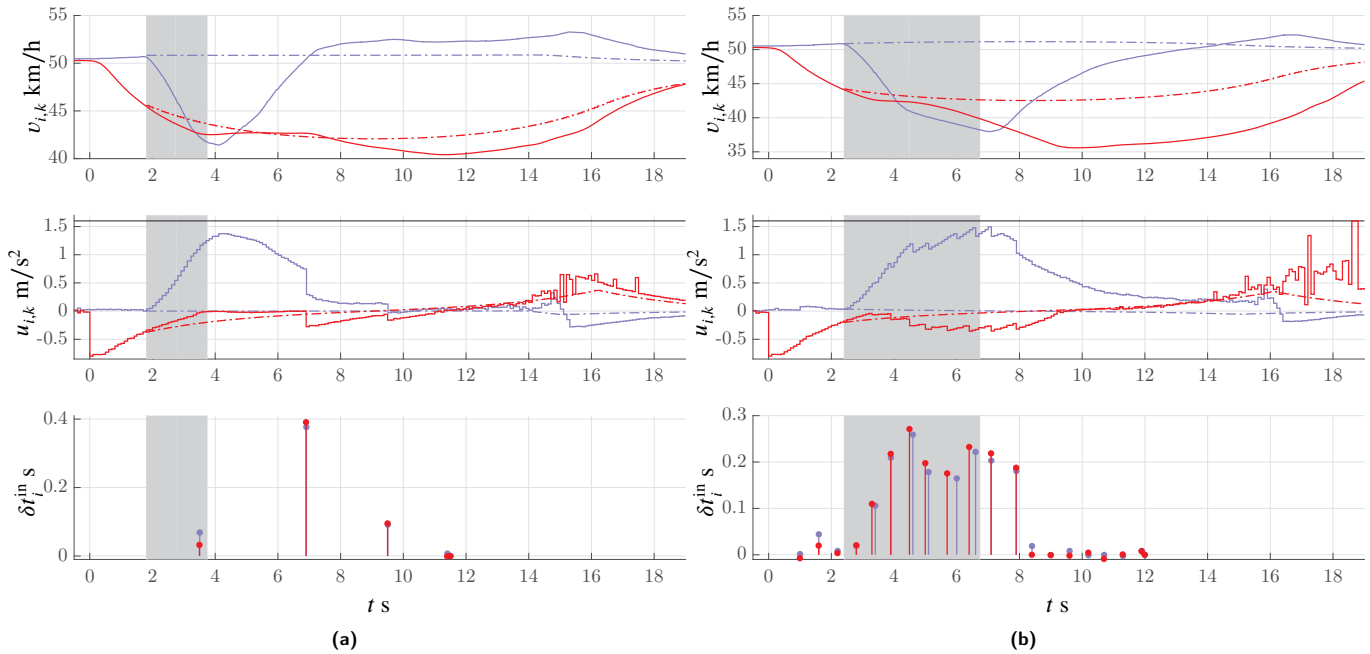
FIGURE 12 Data from experimental scenarios with parameters $Q_1 = 10$, $R_1 = 100$, $Q_2 = Q_3 = 1$,, $R_2 = R_3 = 10$ and $v_i^{\mathrm{ref}} = 50$ km/h where the driver of the first vehicle press the brake. To ease visualization, the data from Vehicle 2 has been removed from the presentation. Subfigure (a) shows the response using the converged intersection-level controller and (b) the response using the 1-step controller. The plots show from top to bottom the velocity, the acceleration commanded by the vehicle-level MPC and the change between consecutive computations of t_i^{in}, δt_i^{in}. In all plots, the gray bar illustrates the time during which the automated system of the first vehicle is overridden by the driver. The dashed lines in the acceleration plots shows the MPC's open loop prediction at the time where the automated system is suspended. The noise in control command of Vehicle 3 around $t \approx 17$ is due to a GPS issue, which is discussed throughly in[23] but omitted here for brevity.

5.2.3 | Rejection of Perturbations

To investigate the ability to counteract large perturbations, experimental scenarios were performed where the drivers of the vehicles overrode vehicle-level controllers for short periods of time by pressing the brake or accelerator pedals. Data from two such experiments where the driver of the first vehicle presses the brake pedal is given in Figure 12 . For comparison, Figure 12 also shows the open-loop predictions made on the onset of the perturbation, which gives an indication of what the unperturbed trajectories would look like.

Rejection of the perturbations should be handled by both feedback loops, i.e. both the action of the vehicles and the time-slot schedule should be adjusted as a consequence. In particular, when the velocity of the first vehicle is reduced due to the perturbation, the time-slot schedule should be adjusted so that the intersection entry of all vehicles are postponed. Indeed, this is also what occurs in both the Converged and 1-step cases, shown in Figure 12 a and Figure 12 b respectively, where the size of the adjustments in t_i^{in} is shown in the lower plots and the perturbation is represented by the gray slab. Note that since the perturbations are introduced manually by the driver, they differ in length and magnitude between Figure 12 a and Figure 12 b .

The benefit of the bi-level control structure is made visible in the middle plots of Figure 12 a. Here it is clearly shown that the application of a recomputed time-slot at $t \approx 7$ reduces the magnitude of the control command of vehicle 1 and increases that of vehicle 2, which effectively distributes the effort required to counter the perturbation among the two vehicles.

A similar behavior, however smaller in magnitude, can be observed for the 1-step controller in Figure 12 b, which causes the jagged behavior in the middle plot. Note that as predicted in Section 3.4, the reaction of the intersection-level control loop is delayed for the converged controller, and the large adjustment to T is not performed until 2.5 s after the perturbation. As discussed in Section 3.4.1, this is due to the usage of a predicted future state as the basis for the solution of the SQP.

A perturbed scenario was also simulated to further highlight the benefits of the bi-level control structure and to enable better comparison between the different intersection-level controllers. The result is presented in Figure 13 , where we also include the
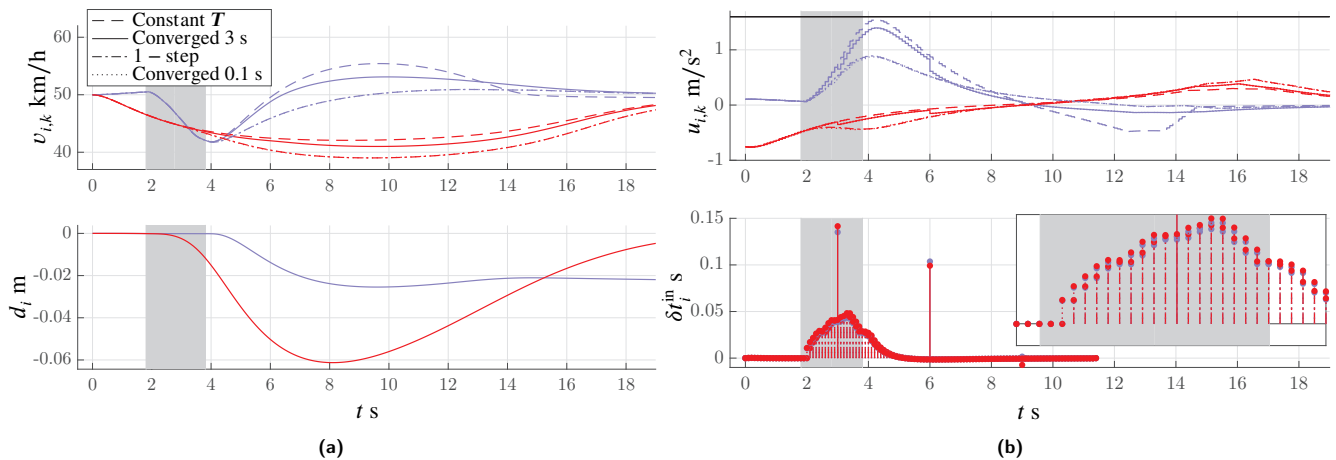
FIGURE 13 Data from simulated scenarios with large perturbations where different versions of the intersection-level controller has been applied. Besides the controllers used in the experiments, the case where the time-slot schedule is kept constant and the case where the SQP is solved to convergence every 0.1 s is shown for comparison. The colors differentiate Vehicle 1 (blue) from Vehicle 3 (red). In the lower plot of subfigure (a), d_i denotes the difference between the position $p_{i,k}$ resulting from the 1-step controller and the controller where the SQP is solved to convergence every 0.1 s. The perturbation is taken from the experimental scenario of Figure 12 a.

case where the time-slots are not adjusted and a highly idealized, unrealistic controller in which the SQP is solved to convergence every $t_s = 0.1$ s at the current state to serve as benchmarks.

Note in particular the trajectories corresponding to the case where the time-slot schedule is kept constant. In this case, the entire effort of perturbation rejection is placed on the first vehicle, with higher transient accelerations and velocities as a consequence. A similar behavior is observed between the time-slot updates when the SQP is solved to convergence every $t_s^T = 3s$. While the effort of rejecting the disturbance is re-distributed among the vehicles with re-computed time-slots, large acceleration levels are observed in Vehicle 1 between $t = 3$ and $t = 6$.

With the 1-step controller, the transient velocities and accelerations are lower as the effort to counter the perturbation is continuously distributed between the vehicles. As can be seen in Figure 13 b, the size of the time-slot adjustments in the 1-step case is very small, and their application results in no rapid changes in the control command. This is due to the higher update frequency, which allows a gradual adjustment of the intersection-level controller to the disturbance.

Finally, the difference between the positions resulting from the the 1-step controller and the idealized case where the SQP solved to convergence every 0.1 s is shown in the lower plot of Figure 13 a. As can be seen, the difference is at most in the order of the accuracy of the positioning system. For most of the time, the difference would not be distinguishable from measurement noise. This is a strong indication that here is no major benefit of solving the SQP to a higher accuracy this fast, further motivating the use of the 1-step controller. We note that the corresponding accumulated difference in t_i^{in} is around 50 ms at most, and settles around 1 ms.

Note that the trajectories for the case where the SQP is solved to convergence every $t = 0.1$ s are drawn in all plots of Figure 13 , but are indistinguishable from the trajectories corresponding to the 1-step controller.

6 | DISCUSSION AND CONCLUSION

This paper addressed the development and experimental validation of a semi-distributed algorithm for optimal coordination of automated vehicles at intersections. In particular, we described a bi-level MPC, where an upper control layer allocates collision free intersection occupancy time-slots by solving a NLP, and a lower control layer computes the optimal control commands by solving QPs. We detailed a semi-distributed SQP method used to solve the NLP, and described a practical implementation of the controller, including the integration of the state-of-the art QP solver HPMPC. We demonstrated the applicability of the controller and computational scheme through an extensive experimental campaign. In particular, even though there were substantial issues

with the implementation, the method was shown to reject large perturbations efficiently and satisfy the collision avoidance constraints to a relevant accuracy.

We want to highlight that the issues observed during the experimental campaign were entirely due to implementation deficiencies and design choices and not inherent to the control formulation or algorithm. On the contrary, the performance observed despite the issues is an indication of the scheme's applicability to real scenarios. For instance, while the ITS-G5 standard has 6 service channels[42], and therefore has the potential to let 6 vehicles communicate simultaneously, larger scenarios would require sequential communication with the current technology. The actual communication time would thereby increase and a delay would be induced: in a scenario with, e.g., 42 cars, at least 7 rounds where 6 vehicles communicate in parallel would be required. With the equipment used during the experiment, the wait on Line 12 of Algorithm 1 would increase with at least 24 ms. Even though a lager scenario likely requires more SQP iterations to converge, as indicated by the results of Section 5.1.2, the time required to solve the problem would still be much smaller than that observed in the experiments, and equal or better performance could be expected. It should also be noted that, even in a non-ideal communication environment where some vehicles possibly need to re-send their data due to packet drops, the added delay is small compared to that experienced during the experiments and would likely not affect the controller performance significantly.

We also want to emphasize the complete parallelizability of the vehicle-level QPs and LPs. Due to this, the time-per-iterate in the SQP will practically be independent of the number of vehicles, and the computational time decided solely by the number of iterates, use of reduced steps, and the time required to solve the QP sub-problems. In terms of computation, the algorithm therefore scales well with an increased number of vehicles.

We should also point out that the central part of the algorithm could be performed at a physically remote location, e.g. in the "cloud". In this case, the coordinating unit would only be required to provide an access point to the V2V network. Moreover, the ability to function with rather large delays could also motivate the use of cellular communication rather than direct radio-links. With a cellular communication solution, no dedicated intersection infrastructure would be needed at all.

As evidenced by both experimental and simulated results, the bi-level controller successfully managed to handle both large and small perturbations. In particular, by closing the upper control loop, we showed that the controller distributed the effort of rejecting perturbations among the involved vehicles. With such a system, the actuation capacities of all involved vehicles can be used to prevent a collision, should it be necessary, which is an important safety feature. Moreover, the 1-step variation of the intersection-level controller was shown to have comparable performance. Indeed, as indicated by the simulation results, the difference between solving the SQP to full convergence and applying the 1-step scheme are on a scale which makes it irrelevant to the application. As the 1-step scheme requires significantly less communication, faster feedback is thereby enabled in the intersection control loop. While this indicates that the 1-step controller is superior to the Converged controller in the tested cases, further studies are required before general conclusions can be drawn.

Future Work

We aim to generalize the method to include scenarios with more than one oncoming lane per road, allow turning vehicles and explicitly account for rear-end collisions between vehicles on the same lane. While rear-end collision avoidance constraints are easy to formulate and include in the centralized setting of Problem (6), they create additional couplings between the vehicles and complicates the application of the decomposition method used in this paper. We are currently working on schemes which allows both rear-end collision avoidance constraints and distributed computation. Finally, we will address the combinatorial problem of finding the optimal crossing order in future work.

ACKNOWLEDGMENTS

References

1. Hult R., Campos G. R., Steinmetz E., Hammarstrand L., Falcone P., Wymeersch H.. Coordination of Cooperative Autonomous Vehicles: Toward safer and more efficient road transportation. *IEEE Signal Processing Magazine.* 2016;33(6):74-84.

2. Colombo A., Del Vecchio D.. Efficient Algorithms for Collision Avoidance at Intersections. *Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control.* 2012;:145–154.

3. Steinmetz E., Hult R., Campos G. R., Wildemeersch M., Falcone P., Wymeersch H.. Communication analysis for centralized intersection crossing coordination. *Proceedings of the 11th International Symposium on Wireless Communications Systems (ISWCS).* 2014;:813-818.

4. Chen L., Englund C.. Cooperative Intersection Management: A Survey. *IEEE Transactions on Intelligent Transportation Systems.* 2016;17(2):570-586.

5. Rios-Torres J., Malikopoulos A. A.. A Survey on the Coordination of Connected and Automated Vehicles at Intersections and Merging at Highway On-Ramps. *IEEE Transactions on Intelligent Transportation Systems.* 2017;18(5):1066-1077.

6. Dresner K., Stone P.. Multiagent traffic management: a reservation-based intersection control mechanism. *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems.* 2004;:530-537.

7. Dresner K., Stone P.. A Multiagent Approach to Autonomous Intersection Management. *Journal of Artificial Intelligence Research.* 2008;31(1):591–656.

8. Kowshik H., Caveney D., Kumar P. R.. Provable Systemwide Safety in Intelligent Intersections. *IEEE Transactions on Vehicular Technology.* 2011;60(3):804-818.

9. Gregoire J., Frazzoli E.. Hybrid centralized/distributed autonomous intersection control: Using a job scheduler as a planner and inheriting its efficiency guarantees. *Proceedings of the 55th IEEE Conference on Decision and Control (CDC).* 2016;:2549-2554.

10. Tallapragada T, Cortés J.. Coordinated intersection traffic management. *IFAC-PapersOnLine.* 2015;48(22):233 - 239. 5th IFAC Workshop on Distributed Estimation and Control in Networked Systems NecSys 2015.

11. Lee J., Park B.. Development and Evaluation of a Cooperative Vehicle Intersection Control Algorithm Under the Connected Vehicles Environment. *IEEE Transactions on Intelligent Transportation Systems.* 2012;13(1):81-90.

12. Campos G. R., Falcone P., Sjöberg J.. Autonomous cooperative driving: A velocity-based negotiation approach for intersection crossing. *Proceedings of the 16th International IEEE Conference on Intelligent Transportation Systems (ITSC).* 2013;:1456-1461.

13. Campos G. R., Falcone P., Wymeersch H., Hult R., Sjöberg J.. Cooperative receding horizon conflict resolution at traffic intersections. *Proceedings of the 53rd IEEE Conference on Decision and Control (CDC).* 2014;:2932-2937.

14. Kim K. D., Kumar P. R.. An MPC-Based Approach to Provable System-Wide Safety and Liveness of Autonomous Ground Traffic. *IEEE Transactions on Automatic Control.* 2014;59(12):3341-3356.

15. Miculescu D., Karaman S.. Polling-systems-based control of high-performance provably-safe autonomous intersections. *Proceedings of the 53rd IEEE Conference on Decision and Control (CDC).* 2014;:1417-1423.

16. Qian X., Gregoire J., La Fortelle A., Moutarde F.. Decentralized model predictive control for smooth coordination of automated vehicles at intersection. *Proceedings of the European Control Conference (ECC).* 2015;:3452-3458.

17. Kamal M. A. S., Imura J., Hayakawa T., Ohata A., Aihara K.. A Vehicle-Intersection Coordination Scheme for Smooth Flows of Traffic Without Using Traffic Lights. *IEEE Transactions on Intelligent Transportation Systems.* 2015;16(3):1136-1147.

18. Murgovski N., Campos G. R., Sjöberg J.. Convex modeling of conflict resolution at traffic intersections. *Proceedings of the 54th IEEE Conference on Decision and Control (CDC).* 2015;:4708-4713.

19. Zhang Y. J., Malikopoulos A. A., Cassandras C. G.. Optimal control and coordination of connected and automated vehicles at urban traffic intersections. *Proceedings of the American Control Conference (ACC).* 2016;:6227-6232.

20. Jiang Y., Zanon M., Hult R., Houska B.. Distributed Algorithm for Optimal Vehicle Coordination at Traffic Intersections. *IFAC-PapersOnLine.* 2017;50(1):11577 - 11582. 20th IFAC World Congress.

21. Katriniok A., Kleibaum P., Joševski M.. Distributed Model Predictive Control for Intersection Automation Using a Parallelized Optimization Approach. *IFAC-PapersOnLine.* 2017;50(1):5940 - 5946. 20th IFAC World Congress.

22. Hult R., Campos G. R., Falcone P., Wymeersch H.. An approximate solution to the optimal coordination problem for autonomous vehicles at intersections. *Proceedings of the American Control Conference (ACC).* 2015;:763-768.

23. Hult R., Zanon M., Gros S., Falcone P.. Optimal Coordination of Automated Vehicles at Intersections: Theory and Experiments `https://arixiv`. 2018.

24. Hult R., Zanon M., Gros S., Falcone P.. Primal Decomposition of the Optimal Coordination of Vehicles at Traffic Intersections. *Proceedings of the 55th IEEE Conference on Decision and Control (CDC).* 2016;:2567-2573.

25. Zanon M, Gros S., Wymeersch H., Falcone P.. An Asynchronous Algorithm for Optimal Vehicle Coordination at Traffic Intersections. *IFAC-PapersOnLine.* 2017;50(1):12008 - 12014. 20th IFAC World Congress.

26. Frison G., Sorensen H.B., Dammann B., Jorgensen J.B.. High-performance small-scale solvers for linear Model Predictive Control. *Proceedings of the European Control Conference (ECC).* 2014;:128-133.

27. Diehl M.. *Real-Time Optimization for Large Scale Nonlinear Processes* Fortschr.-Ber. VDI Reihe 8, Meß-, Steuerungs- und Regelungstechnik, vol. 920: . Düsseldorf: VDI Verlag; 2002. Download also at: http://www.ub.uni-heidelberg.de/archiv/1659/.

28. Nocedal J., Wright S.J.. *Numerical Optimization.* Springer Series in Operations Research and Financial EngineeringSpringer; 2 ed.2006.

29. Büskens C., Maurer H.. Online Optimization of Large Scale Systemsch. Sensitivity Analysis and Real-Time Optimization of Parametric Nonlinear Programming Problems, :3–16. Berlin, Heidelberg: Springer Berlin Heidelberg 2001.

30. Diehl M., Findeisen R., Schwarzkopf S., et al. An Efficient Algorithm for Nonlinear Model Predictive Control of Large-Scale Systems. Part I: Description of the Method. *Automatisierungstechnik.* 2002;50(12):557–567.

31. Findeisen R., Allgöwer F.. Computational Delay in Nonlinear Model Predictive Control Proc. Int. Symp. Adv. Control of Chemical Processes, ADCHEM2003.

32. Chen W., Ballance D.J., O'Reilly J.. Model Predictive Control of Nonlinear Systems: Computational Delay and Stability. *IEEE Transactions on Automatic Control.* 2000;147(4):387–394.

33. Fletcher R.. *Practical Methods of Optimization.* Chichester: Wiley; 2nd ed.1987.

34. Axehill D.. Controlling the level of sparsity in MPC. *Systems & Control Letters.* 2015;.

35. Diehl M., Ferreau H. J., Haverbeke N.. Nonlinear model predictive control Lecture Notes in Control and Information Sciences, vol. 384: ch. Efficient Numerical Methods for Nonlinear MPC and Moving Horizon Estimation, :391–417. Springer 2009.

36. Hult R., Sancar E., Jalalmaab M., et al. Design and experimental validation of a cooperative driving control architecture for the Grand Cooperative Driving Challenge 2016. *Transactions on Intelligent Transportation Systems.* 2017;. to Appear.

37. RENDITS `http://www.rendits.com/`Accessed: 2016-11-05; .

38. Zanon M., Gros S., Diehl M.. A Tracking MPC Formulation that is Locally Equivalent to Economic MPC. *Journal of Process Control.* 2016;.

39. Hult R., Zanon M., Gros S., Falcone P.. Optimal coordination of three cars approaching an intersection `https://youtu.be/nYSXvnaNRK4`Accessed: 2017-03-03; .

40. Fernandez J. A., Borries K., Cheng L., Kumar B. V. K. Vijaya, Stancil D. D., Bai F.. Performance of the 802.11p Physical Layer in Vehicle-to-Vehicle Environments. *IEEE Transactions on Vehicular Technology.* 2012;61(1):3-14.

41. Rittger L., Schmidt G., Maag C., Kiesel A.. Driving behaviour at traffic light intersections. *Cognition, Technology & Work.* 2015;17(4):593–605.

42. European Telecommunications Standards Institute . *ETSI EN 302 663: Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band.* 2012.