Transportation Research Part C 62 (2016) 117–132

Contents lists available at ScienceDirect

Transportation Research Part C

journal homepage: www.elsevier.com/locate/trc

Hierarchical model predictive control for multi-lane motorways in presence of Vehicle Automation and Communication Systems

Claudio Roncoli [*], Ioannis Papamichail, Markos Papageorgiou

Dynamic Systems and Simulation Laboratory, Technical University of Crete, Chania 73100, Greece

A R T I C L E   I N F O

Article history:
Received 9 April 2015
Received in revised form 16 November 2015
Accepted 16 November 2015
Available online 30 November 2015

Keywords:
Motorway traffic control
Vehicle automation and communication systems
Model predictive control

A B S T R A C T

A widespread deployment of vehicle automation and communication systems (VACS) is expected in the next years. This may lead to improvements in traffic management efficiency because of the novel possibilities of using VACS both as sensors and as actuators, as well as of a variety of new communications channels (vehicle-to-vehicles, vehicle-to-infrastructure) and related opportunities. To achieve this traffic flow efficiency, appropriate studies, developing potential control strategies to exploit the VACS availability, are essential. This paper describes a hierarchical model predictive control framework that can be used for the coordinated and integrated control of a motorway system, considering that an amount of vehicles are equipped with specific VACS. The concept employs and exploits the synergistic (integrated) action of a number of old and new control measures, including ramp metering, vehicle speed control, and lane changing control at a macroscopic level. The effectiveness and the computational feasibility of the proposed approach are demonstrated via microscopic simulation for a variety of penetration rates of equipped vehicles.

1. Introduction

The problem of traffic congestion in and around densely inhabited areas has a strong economical and social impact. One possible solution is the construction of wider road infrastructures, with an enormous economical cost and significant environmental consequences. On the other hand, the currently existing motorways are actually underutilised, especially in the periods of high demand due to congestion (Papageorgiou et al., 2003). A possible way to overcome this situation is the development and implementation of proper traffic control measures and strategies with the aim of reducing traffic congestion and increasing the overall capacity of traffic networks.

In the last two decades, a significant and increasing steadily interdisciplinary effort by the automotive industry, as well as by numerous research institutions around the world, has been devoted to planning, developing, testing and deploying a variety of Vehicle Automation and Communication Systems (VACS) that are expected to revolutionise the features and capabilities of individual vehicles within the next decades (Bishop, 2005). Among the wide range of proposed VACS, only few have actually a direct impact on traffic flow, since the majority of VACS aims at primarily improving safety or driver convenience (Diakaki et al., 2015). Some VACS may thus be exploited to interfere with the driving behaviour via recommending, supporting, or even executing appropriately designed traffic control tasks. This gives the possibility of having access to control actions that are not available with conventionally driven cars (e.g., individual vehicle speed or lane-change advice). On the other hand, the uncertainty in the future development of VACS calls for the design of control strategies that are robust with respect to the possible types of VACS, as well as to their penetration rate.

---

[*] Corresponding author.
E-mail addresses: croncoli@dssl.tuc.gr (C. Roncoli), ipapa@dssl.tuc.gr (I. Papamichail), markos@dssl.tuc.gr (M. Papageorgiou).

http://dx.doi.org/10.1016/j.trc.2015.11.008
0968-090X/© 2015 Elsevier Ltd. All rights reserved.

The use of an intelligent and connected infrastructure for traffic management has been considered in the Automated Highway System (AHS) concept (Varaiya, 1993), where it was assumed that platoons of fully automated vehicles may travel on specifically designed motorways. This complex system was suggested to be controlled via a multi-layer control structure, where the traffic-level control strategies are included in a decentralised link-layer. One of the first works addressing link-layer control strategies was proposed by Rao and Varaiya (1994). More recently, Baskar et al. (2012) proposed a model predictive control (MPC) approach for the integrated control (addressing speed, lane assignment and ramp metering) of platoon-based AHS, that involves both real-valued and integer variables, leading to a mixed non-convex optimisation problem that may be difficult to solve in real-time. A number of other works addressed specifically the problem of deciding on efficient vehicle lane-paths for a motorway under fully automated (AHS) or semi-automated driving (e.g. Hall and Lotspeich, 1996; Ramaswamy et al., 1997; Kim et al., 2008). However, to tackle the problem complexity, a number of simplifying assumptions were typically made, such as known and constant prevailing speeds along the highway and absence of traffic congestion, thanks to the assumed (but not addressed) operation of ramp metering (RM) at the highway entrances; also, a number of structural assumptions were made to limit the (otherwise vast) space of potential lane-path assignments.

On the other hand, the coordinated and integrated exploitation of conventional traffic control actuators, such as road-side traffic signals and variable message signs (VMS) for route guidance, variable speed limits (VSL), and RM, has been proposed in several papers. Some approaches are based on the formulation of appropriate optimisation problems, envisioning their application within an MPC scheme (Kotsialos et al., 2002; Hegyi et al., 2005a; Gomes and Horowitz, 2006; Papamichail et al., 2010; Frejo et al., 2014; Chow, 2015; Ferrara et al., 2015). Nevertheless, the intrinsic complexity of these approaches may be an impediment for real-time application while also considering additional options and features offered by emerging VACS. Additional difficulties may appear due to the non-convex nature of the related optimal control problems. Other control strategies where designed following an analysis of some properties of traffic dynamics, e.g. Hegyi et al. (2005b), Sun and Horowitz (2005), Zhang et al. (2006), Hegyi et al. (2008), Muralidharan and Horowitz (2012), Torné et al. (2014).

The purpose of this paper, that represents an extension of Roncoli et al. (2014), is the development of a hierarchical control framework based on an MPC scheme for the coordinated and integrated motorway traffic management, taking into account the possibility of using VACS both as sensors and as actuators, with the advantages of having an increased degree of freedom with respect to the control possibilities, as well as a more precise estimation of the motorway state, compared to conventional systems. In particular, according to the nomenclature on automated motorway traffic control (see, e.g., Varaiya, 1993), this paper deals with the so-called "link layer", aiming at smoothing and improving traffic conditions. Therefore, problems at higher levels (e.g., route assignment) or at lower levels (e.g., car-following laws) are assumed to be properly addressed by other (external) systems. It is supposed that VACS-equipped vehicles have the capability of bidirectional communication with the infrastructure (V2I); appropriate control actions are decided in a centralised manner by a Traffic Management Center (TMC) and dispatched to specific vehicles for their implementation. In Fig. 1, the envisioned scenario is sketched. The core of the methodology is the convex optimisation problem proposed by Roncoli et al. (2015b), that is based on the piecewise linear macroscopic traffic flow model introduced by Roncoli et al. (2015a), which considers, as decision variables, actions that are enabled with the aid of VACS. Since the application of this methodology in a real motorway environment will not be possible for several years to come, because of the necessary amount of vehicles equipped with appropriate devices, the best opportunity to realistically test the proposed control strategy is by use of a microscopic traffic simulator; this latter aspect is the main issue considered in this paper.

The paper is structured as follows: Section 2 describes the proposed control framework. In Section 3, the microscopic simulation environment is described, while in Section 4 the obtained simulation results are presented and compared with a reference no-control case. Section 5 concludes the paper, highlighting the main results and introducing some challenging research tasks for the future.

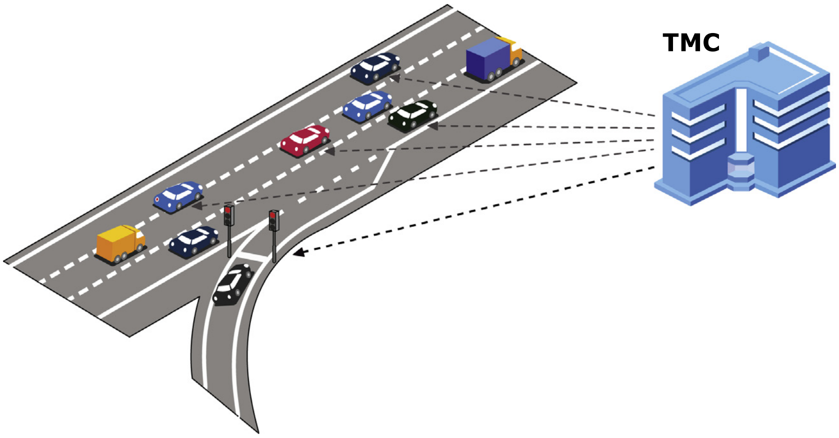

Fig. 1. The envisioned scenario, where vehicles communicate with a TMC, that computes and dispatches control actions.

2. Control framework

2.1. The control structure

Motorway traffic flow, like many other complex processes, is affected by several factors and any related mathematical model has necessarily a limited accuracy. On the other hand, the employed model must be simple enough to allow for computational tractability of the related optimal control problem. Thus, the use of an open-loop control strategy (whereby the control trajectories are computed at the initial instant, without being updated during the process) may lead to increasingly diverging process behaviour, compared with the predicted one, due to inaccuracies in predicting the external disturbances (mainly the demands) or limited model accuracy. A mitigation of these issues is offered by the utilisation of a receding horizon (or MPC) scheme, that entails that the control actions are re-computed periodically, using updated measurements and predictions (Camacho and Bordons, 1995). This permits to reject past inaccuracies and to maintain the difference between the model predictions and the real process outcome at low levels, thus improving the overall control performance. For operational feasibility and maximum efficiency, the MPC problem may be cast in a multi-layer control structure sketched in Fig. 2.

The related optimal control problem for the present application is included in the optimisation layer and is detailed in Section 2.3. The information retrieved from fixed sensors or moving vehicles cannot be fed directly to the optimisation layer, as it requires to be previously appropriately filtered, aggregated, and processed. In the proposed framework, these tasks are performed within the adaptation and prediction layer (Section 2.2), that has in fact the purpose to compute the current traffic state and the estimated demand for a given optimisation horizon. Moreover, the output of the optimisation layer is not directly ready for implementation in the motorway network, therefore a further application layer (Section 2.5) is introduced, aiming at converting the macroscopic optimal results into applicable control tasks sent to actuators, which may be infrastructure-based (e.g. traffic lights in the RM case) or installed within vehicles (with V2I capabilities). In order to account for model inaccuracies or unpredicted local perturbations, also a local control layer may be introduced (Section 2.4), which modifies the optimal results (provided by the optimisation layer) to account for local short-term variations.

It should be noted that the multi-layer control structure developed follows the principles and concept proposed in the pioneering work by Findeisen et al. (1980), which has found numerous applications in various domains. Within the traffic

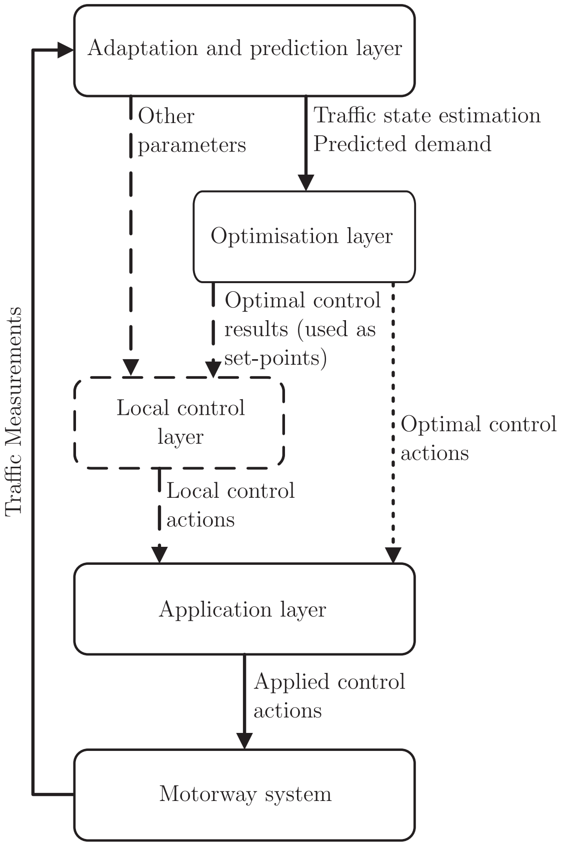

Fig. 2. The proposed multi-layer control structure.

context, similar control structures were also proposed by Papageorgiou (1984) and Papamichail et al. (2010). In the following sections, each layer is described in more detail.

2.2. The adaptation and prediction layer

The purpose of this layer is essentially to process the data retrieved from the motorway system in order to obtain necessary information to be used by the lower layers; as previously mentioned, the main performed tasks are the estimation of the current traffic state and the prediction of future demands.

With conventional vehicles, all the data available are retrieved from road-side traffic sensors, that are placed at specific locations of the motorway (sometimes quite distant from each other). The presence of VACS with V2I capability provides the opportunity to extend and enhance the real-time measurement capabilities via available vehicle information from on-board sensors, such as vehicle speed, position, and distance to the surrounding vehicles. These data may be shared with other vehicles (V2V communication) or with the infrastructure (V2I communication). These new possibilities lead to an unprecedented accuracy, richness, and granularity of available real-time information which opens new avenues for real-time estimation of the traffic state. Related traffic state estimation approaches have been reported, e.g. by de Fabritiis et al. (2008), Van Lint and Hoogendoorn (2010), Herrera and Bayen (2010), Yuan et al. (2012), Seo et al. (2014), Bekiaris-Liberis et al. (2015).

Traffic demand estimation is a complex task of crucial importance in an MPC framework, since the results of the optimisation problem are strongly influenced by a proper forecasting of the demand expected during the defined optimisation horizon; classic forecasting models are based on measurements and historical data elaboration (e.g. Zhou and Mahmassani, 2007). Again, a high penetration of VACS may give the possibility of improving the knowledge on the number of vehicles that are approaching a specific area, permitting to improve the demand prediction accuracy.

2.3. The optimisation layer

The optimisation layer contains the numerical solution of an optimisation problem, which is solved periodically, at predefined control intervals in the order of minutes. Since the numerical solution is computed in real-time, a crucial aspect is the time needed to obtain it. In the present approach, this was the motivation for the definition of a convex Quadratic Programming (QP) problem, which can be solved very efficiently with available algorithms. The traffic modelling aspects considered in this model have been described in detail by Roncoli et al. (2015a,b); hereafter a brief account of the modelling aspects is provided for self-completeness.

The given motorway stretch is subdivided into segments (indexed by $i = 1, \ldots, I$) and lanes (indexed by $j = 1, \ldots, J$); considering a discrete time step T and a given optimisation horizon K, the discrete time index of the employed model is $k = 1, \ldots, K$, where $t = kT$. The following three control actions, each one characterised by a specific (different) control time step, may be activated:

- Ramp-metering (RM): It regulates the inflow from the on-ramps to the motorway mainstream and is currently applied on many motorways (see, e.g. Papamichail et al., 2010); the corresponding control variable $r_{i,j}(k^R)$ [veh/h] denotes the controlled ramp outflow, where the ramp is located at segment i, lane j, during control interval $(k^R, k^R + 1]$, where $k^R = \left\lceil \frac{kT}{T^R} \right\rceil$, and T^R is the control step for RM.
- Mainstream Traffic Flow Control (MTFC) via VSL: The use of VSL for traffic management has been exploited in an increasing number of research works; among others Smulders and Helleman (1998), Breton et al. (2002), Hegyi et al. (2005b), Carlson et al. (2010), Nissan and Koutsopoulos (2011), Yang et al. (2013), Chen et al. (2014). In this work, a different VSL may be imposed for each segment-lane of the motorway by ordering the speed of VACS-equipped vehicles, aiming finally at achieving a specific (optimal) mainstream flow value; the longitudinal optimal flows for segment i, lane j are defined by $q_{i,j}(k^Q)$ [veh/h], where $k^Q = \left\lceil \frac{kT}{T^Q} \right\rceil$ and T^Q is the control step for MTFC.
- Lane Changing Control (LCC): The optimal lateral flows are computed for each segment-lane, thus enabling an optimal distribution of traffic flow among the different lanes; lateral flows are represented by variables $f_{i,j,\bar{j}}(k^F)$ [veh/h], that denote the amount of vehicles moving from lane j to lane $\bar{j}$ ($\bar{j} = j \pm 1$), remaining within the same segment i, during control interval $(k^F, k^F + 1]$, where $k^F = \left\lceil \frac{kT}{T^F} \right\rceil$ and T^F is the control step for LCC.

The dynamic equation for densities $\rho_{i,j}$ [veh/km] for each segment-lane (i,j), is given by the following conservation equation:

$$\rho_{i,j}(k+1) = \rho_{i,j}(k) + \frac{T}{L_i}\left[q_{i-1,j}(k^Q) + r_{i,j}(k^R) - q_{i,j}(k^Q) - \gamma_{i,j}(k)\sum_{j=1}^{J} q_{i-1,j}(k^Q) + f_{i,j+1,j}(k^F) + f_{i,j-1,j}(k^F) - f_{i,j,j-1}(k^F) - f_{i,j,j+1}(k^F)\right] \tag{1}$$

where $\gamma_{i,j}(k)$ are estimated turning rates at off-ramps. A graphical representation of the variables related to each segment-lane entity is provided in Fig. 3. The possibility of performing RM actions may lead to the creation of queues $w_{i,j}$ [veh] at on-ramps, whose dynamics are described by the following conservation equation:

$$w_{i,j}(k+1) = w_{i,j}(k) + T\left[d_{i,j}(k) - r_{i,j}(k^R)\right] \tag{2}$$

where $d_{i,j}(k)$ [veh/h] is the external (predicted) demand feeding the model.

As mentioned earlier, longitudinal flows are considered as control variables that can be realised via appropriate VSL actions. Therefore, they are constrained from above by the longitudinal flow values that would prevail without VSL, i.e. by their uncontrolled values, which are obtained according to a discretised first-order traffic flow model with piecewise linear fundamental diagram (FD) based on a modified Godunov-discretised LWR model (Lighthill and Whitham, 1955; Richards, 1956) that, differently from the classic CTM (Daganzo, 1994), includes specific terms to account for the capacity drop phenomenon (see Roncoli et al., 2015a for details). It is well-known that the Godunov-discretised longitudinal flow of first-order LWR-based (Lighthill and Whitham, 1955) traffic flow models may be viewed as the minimum between a demand function, which depends on the upstream-segment density, and a supply function, which depends on the downstream-segment density. In the present formulation, the capacity drop is introduced via a modification of the demand function; specifically, if the upstream density is over-critical, the demand function is not constant and equal to capacity (as in the ordinary Godunov-discretised model), but is linearly decreased in dependence of the density, according to an approach first proposed by Lebacque (2003), see Fig. 4; also, the entering flow from an on-ramp (if any) and from the adjacent segments (lateral flows) are modelled to decrease capacity, since the related acceleration or braking of vehicles usually perturb the mainstream traffic flow.

Other approaches defining optimal control problems, albeit for conventional traffic flow, including CTM-based models can be found in the work by Gomes and Horowitz (2006), where the authors proposed a linear program (LP) for optimal RM; Li et al. (2014) proposed a centralised optimal control problem formulated as a mixed-integer linear program (MILP) for RM, VSL, and hard-shoulder running, accounting explicitly for different FDs that materialise while applying different speed limits; while Chow (2015), extending the work by Li et al. (2014), proposed also parsimonious decentralised control strategies based upon aggregated traffic information. The main differences of these works from the employed model are: (i) the introduction of lateral flows as control variables that are subject to physical constraints; (ii) the introduction of capacity drop; (iii) the possibility of employing VSL by lane and at arbitrary space locations (no gantries).

As mentioned earlier, the optimisation problem for coordinated and integrated motorway traffic control in presence of VACS has a convex QP form. Specifically, the quadratic cost function includes (see Roncoli et al., 2015a for a detailed explanation):

- Linear terms reflecting the Total Time Spent (TTS), as the most crucial control objective, which includes both the total travel time on the mainstream and the total waiting time at on-ramp queues, and a penalty term to avoid excessive lateral (lane-changing) flows.
- Quadratic penalty terms to reduce time variations of RM and LCC control variables, as well as to reduce time and space fluctuations of the speed values (approximated via appropriate linearised expressions).

Appropriate weights are introduced for each term to reflect the respective control priorities; in particular, the weights related to lateral flows may be tuned in order to encourage lane-changes at specific segment-lanes, e.g. upstream of on-ramps or lane-drops.

In summary, the problem is described via densities and on-ramp queues as state variables; ramp flows, mainstream (longitudinal) flows, and lateral (lane-changing) flows as the control variables. The dynamics (linear equality constraints) comprise the linear conservation Eqs. (1) and (2); while linear inequality constraints take into account the piecewise-linear terms related to longitudinal and lateral flows in the form of upper-bounds for the respective control variables. Fixed upper bounds (capacities) are also considered for the on-ramp queues and flows, and for the off-ramp flows. Finally, non-negativity constraints are specified for all the variables. A detailed description of the developed QP problem (with the quadratic cost function and all linear equality and inequality constraints) may be found in the original paper by Roncoli et al. (2015b).

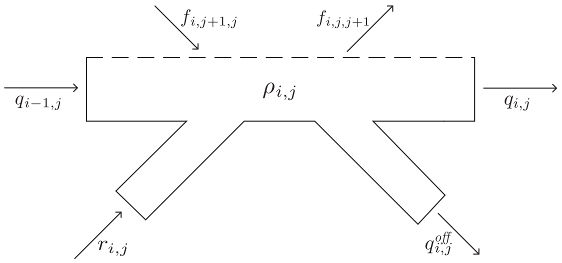

Fig. 3. The segment-lane variables used in the model formulation.

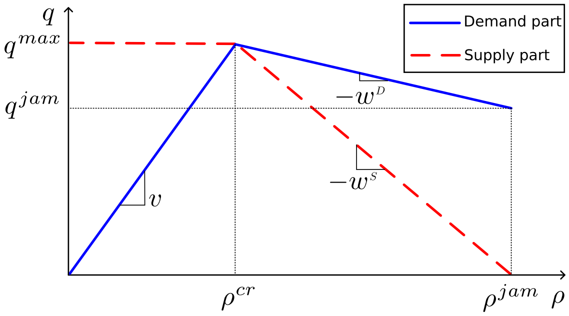

Fig. 4. The demand and supply functions of the used fundamental diagram: the flow q leaving a link is constrained, in congested state ($\rho > \rho^{cr}$), by a linearly decreasing line (with slope $-w^D$), generating a reduction of capacity in congestion, mimicking the capacity drop phenomenon.

2.4. The local control layer

A direct application of the computed optimal control actions may provide an improvement of traffic conditions, and, thanks to the feedback nature of the MPC approach, reject to some extent prediction errors and model inaccuracies. On the other hand, the model used in the optimisation layer includes some simplifying assumptions; for example, while the longitudinal flow of each segment-lane is assumed controllable, the real control inputs are VSL; also, the update time step for the optimisation problem can be in the order of minutes, leading to non-negligible deviation of the real traffic conditions from their computed optimal values. In addition, if the available storage space of an on-ramp is about to be exhausted, specific queue management algorithms (e.g., Smaragdis and Papageorgiou, 2003) may increase the outflow of the ramp to higher values than the optimal ones. For these reasons, further improvements may be obtained by introducing a local control layer (Fig. 2), which incorporates a series of simple decentralised and fast (e.g. 20-s sampling) feedback loops.

The basic idea for the local control layer is to use some values resulting from the optimisation problem as set-points for simpler decentralised (local) feedback controllers, which have the purpose of ensuring that the computed optimal state trajectories are actually implemented with increased accuracy. Because of their major impact on traffic flow evolution, the local controllers act on RM and MTFC via VSL; while the lane-change decisions are left unaltered, as computed by the optimisation layer. The optimal densities related to each segment-lane are chosen to be used as set-points for the local layer controllers for two main reasons: first, they have a direct impact on the TTS, whose minimisation is the main target of this work; and, second, because they are generally characterised by lower fluctuations than equivalent flows.

Regarding the VSL, which apply to individual segment-lanes, the proposed approach implies that the regulator set-point is related to the optimal density of a specific segment-lane; this could be the same segment of VSL application or the downstream segment. We recall that the main function of VSL is to control the mainstream flow control upstream of a potential bottleneck, e.g. in cases of a lane drop or an uphill (sag) (Goñi Ros et al., 2014) or tunnel or queue management at a full on-ramp (Carlson et al., 2014). To this end, it is necessary for the VSL feedback regulators to target the optimal density of the next downstream segment-lane, so as to reduce the arriving flow, in case a density higher than the set-point is measured. A simple I-type controller is suitable for this task, as also discussed by Muller et al. (2015). The following equation describes the dynamics of the employed I-type controller:

$$v_{i,j}(k) = v_{i,j}(k-1) + K_v[\hat{\rho}_{i+1,j}(k) - \bar{\rho}_{i+1,j}(k)] \tag{3}$$

where

- $v_{i,j}(k)$ is the speed limit applied at segment-lane (i,j);
- $\hat{\rho}_{i,j}(k)$ is the density set-point of segment-lane (i,j);
- $\bar{\rho}_{i,j}(k)$ is the measured density of segment-lane (i,j);
- K_v is the integral gain.

As mentioned earlier, the set-point is normally set as the density resulting from the optimisation problem; however, in case the predicted demand (or the traffic behaviour) is not exact (e.g., the demand is underestimated), it may happen that the density set-point assumes an unrealistic low value, calling for unnecessary reduction of speed limit. To overcome this problem, the set-point is set equal to the critical density ($\hat{\rho}_{i,j}(k) = \rho^{cr}_{i,j}$) in case the condition $\hat{\rho}_{i,j}(k) < \bar{\rho}_{i,j}(k) < \rho^{cr}_{i,j}$ is met.

Some rules are eventually applied to the computed $v_{i,j}(k)$, e.g. avoiding that the difference between successive speeds (in time and/or in space) is larger than a threshold, or setting lower- and upper-bounds to the posted VSL. A detailed description of potentially useful rules is proposed by Carlson et al. (2011). Moreover, it has been observed that the capacity flow vs. speed limit relation is actually characterised by a non-linear behaviour (Muller et al., 2015); this makes a linear control inadequate to maintain good damping properties at all allowed speed limits. To circumvent this problem, gain scheduling

(Bishop and Wittenmark, 1995) is adopted, where different integral gains are assigned for different operation points. The current measured speed can be used to determine the current operation point, considering then a look-up table to select an appropriate gain.

Analogously, also RM can be addressed by the local control layer using an I-type controller, as applied within the well-known ALINEA control strategy (Papageorgiou et al., 1991). In order to avoid interferences during the application of the two control tasks, VSL is applied in all segments except from the ones immediately upstream of each merging segment, where the following controller is applied:

$$r_{i,j}(k) = r_{i,j}(k-1) + K_r\left[\hat{\rho}_{i,j}(k^*) - \bar{\rho}_{i,j}(k)\right]$$

where

- $\hat{\rho}_{i,j}(k^*)$ is the density set-point of segment-lane (i,j), lower-bounded by the critical density;
- $\bar{\rho}_{i,j}(k)$ is the density measured at the merge segment-lane (i,j);
- K_r is the integral gain.

The resulting ramp-flow values are appropriately bounded according to minimum and maximum ramp flows allowed.

2.5. The application layer

The application layer is needed in order to convert the outcome of the optimisation or local control layers to actual control actions to be applied in the motorway system. Specifically, it includes procedures for handling the three addressed control actions.

The application of RM actions can be performed using ordinary traffic signals at on-ramps, via appropriate green and red phases, which depend on the computed ramp outflows, as detailed by Papageorgiou and Papamichail (2008). Alternatively, in presence of VACS, the same impact can be obtained providing the commands directly through an in-car information system. When RM is applied, the created ramp queues may exceed the corresponding allowed upper limits. Creation of excessive ramp queues can be avoided with the application of a queue control policy (Smaragdis and Papageorgiou, 2003).

The application of VSL can be improved by an appropriate use of VACS. In fact, in the conventional case of manually driven vehicles, the application of a speed limit is effectuated by the use of VMS located on gantries which display the same speed limit for all lanes. The granularity of these actions is also dependent on the distance between successive VMS and cannot be changed in real time. The use of VACS may drastically upgrade the possibilities of applying VSL. In fact, supposing that a sufficiently high number of vehicles is equipped with V2I communication, each equipped vehicle can receive a specific speed limit (or a suggested cruise speed) that should be respected while driving in the current location. In this case, the spatial granularity of the action is completely customisable by the control system, permitting to arbitrarily modify the application areas and lanes without expensive modifications of the infrastructure. A possible further step in this direction could be the integration within Adaptive Cruise Control (ACC) or Cooperative ACC (VanderWerf et al., 2001), setting the desired speed directly in the vehicle driving systems, without requiring any intervention by the driver. It should be noted that a sufficient penetration of equipped vehicles will be effective to impose the speed limit to non-equipped vehicles as well.

The implementation of LCC actions is more cumbersome, even if all vehicles are in communication with the control center. The control actions can be implemented by sending lane-changing advices to an appropriate number of selected vehicles; the selection may be based on the known destinations of the vehicles and further criteria. Since, for a foreseeable future, the lane change advice will not be mandatory, the assignment will have to account for the compliance rate, as well as for other, spontaneous lane-changes decided by the drivers; the latter may be reduced by involving additional “keep-lane” advices to all equipped vehicles that do not receive a lane-change advice. Cooperative lane-changing possibilities of vehicles equipped with V2V communication capabilities may further facilitate the LCC action. Clearly, any mismatch between the optimal lateral flows and the actually triggered lane changes may be partially compensated thanks to the feedback included in the optimisation layer (MPC).

While more complex cases are subject of ongoing work, the following experiments in this paper are based on the assumptions that equipped-vehicles on the mainstream can receive a lane-changing advice and that the concerned drivers promptly follow these instructions, subject only to physical constraints that may disallow a vehicle to actually change its lane.

3. Experimental setup

3.1. Microscopic simulator setup

The proposed control methodology has been implemented and tested within the microscopic traffic simulator AIMSUN (TSS – Transport Simulation Systems, 2014). The standard configuration of this tool is based on car-following and lane-changing behavioural models derived from the Gipps Model (Gipps, 1981, 1986). However, these models have been reported to have two considerable drawbacks: first, the Gipps car-following model is often not reproducing a realistic capacity drop at the head of congestion (Wang et al., 2005); second, the ability to accurately capture the merging behaviour in a critical flow regime has been criticised (Chevallier and Leclercq, 2009).

In order to overcome the first issue, the Gipps car-following model has been replaced here with the Intelligent Driver Model (IDM) (Treiber et al., 2000), as applied by Ntousakis et al. (2015), that is deemed to provide more realistic results while reproducing the capacity drop. In addition, IDM is deemed to be capable of reproducing both manual driver behaviour and ACC-equipped vehicles (Kesting et al., 2007); therefore the defined traffic behaviour, although not explicitly required, can be considered also as representative for a mixed-scenario with conventional and VACS-equipped vehicles.

The second issue has been tackled with the introduction of some heuristic rules that override the AIMSUN lane-changing policies, specifically at merge areas, so as to obtain more realistic merging situations. Note that, in AIMSUN (as in most real infrastructures as well), an on-ramp leads to an acceleration lane which runs parallel to the mainstream lanes for some 200 m. Thus, vehicles exiting the on-ramp enter the acceleration lane and need eventually to change lane in order to enter the mainstream before the acceleration lane drop. In the modified lane-changing model, a vehicle is allowed to switch to the mainstream if a number of defined conditions are satisfied; they include the vehicle position on the acceleration lane, its current speed, the relative speed with respect to the mainstream vehicles, and the available gap in the mainstream (the gap is calculated as a function of the space, the speed of the merging vehicle, and the speeds of the upstream and downstream mainstream vehicles travelling within the target lane). Similar rules are also applied to all other lanes of the merge area, albeit using different parameter values than for the shoulder lane. The modified model was visually observed to produce a realistic merging behaviour under many different scenarios and flow levels. Nevertheless, calibrating the modified merging model with real data is an interesting, though non-trivial task, which requires a high amount of real microscopic data; and is certainly beyond the scope of the present work, which focuses on testing and evaluating a comprehensive motorway traffic control strategy.

In the performed experiments, equipped vehicles are applying the control actions whenever it is possible (i.e., in case their application does not violate any physical constraint). For MTFC, whenever an equipped vehicle enters a segment-lane, its maximum speed is set according to the value computed by the controller and acts as an upper-bound within its car-following rules.

In the LCC case, all spontaneous lane-changes are inhibited for equipped vehicles; instead, the following logic is applied: each VACS-equipped vehicle computes the available gaps for the left and right lanes using on-board sensors (e.g., in reality, radar sensors) and continuously communicates this information to the TMC; at each control step, the TMC computes the optimal lateral flow for each segment-lane, then converts the obtained value into time intervals between two consecutive lane-changes; at the end of each determined time-interval, a change-lane command is sent to the vehicle characterised by the highest gap within the corresponding segment-lane, which implies that it is the vehicle that will produce the less negative impact on traffic behaviour after changing lane. In case no vehicles are capable to change lane as requested, the last command is repeated within the next simulation step(s). An appropriate simple procedure is also applied to handle the transition between successive time intervals between two consecutive lane-changes.

Considering the suggested (and utilised) control steps, which are in the order of 20 s–1 min, and the proposed size of segments, which are in the order of 500 m, sensor delays (commonly around 0.1–0.2 s) and communication delays (commonly around 0.1–0.4 s) are not considered, as they will not have a major impact on the produced results.

3.2. Network description

A set of experiments is performed on a motorway stretch of 5 km in length, composed of three lanes ($j = 1, 2, 3$ from the shoulder lane to the median lane), with an on-ramp placed at 3.5 km from the motorway entrance. The on-ramp leads to an acceleration lane of 190 m in length. Traffic signals for RM are placed at 20 m upstream of the acceleration lane entrance. The network is composed of 10 homogeneous sections of 500 m in length, whereby the on-ramp is placed in segment 8, as shown in Fig. 5. The use of detectors is not necessary, since it is assumed that the traffic state can be obtained sufficiently accurate via V2I communication systems (see e.g. Bekiaris-Liberis et al., 2015).

The entire simulation horizon is 50 min. The average traffic demand is set according to a trapezoidal shape (see Fig. 6); however, assuming that the vehicle arrivals follow a Poisson distribution, the time intervals between two consecutive arrivals (headway) are sampled according to an exponential distribution. This permits to have different actual demand profiles for different simulation replications, which may be deemed to emulate a recurring traffic pattern that may appear at the same location on different days.

4. Results

4.1. Performed simulations

A broad set of simulations has been performed to calibrate and evaluate the proposed control strategies. First, the macroscopic model, the optimisation problem, and the local control layer are calibrated and tuned using one single replication; then the obtained results are validated using some simulation replications; subsequently different scenarios characterised by different assumptions related to different control strategies are simulated as reported in the next sections. Each scenario

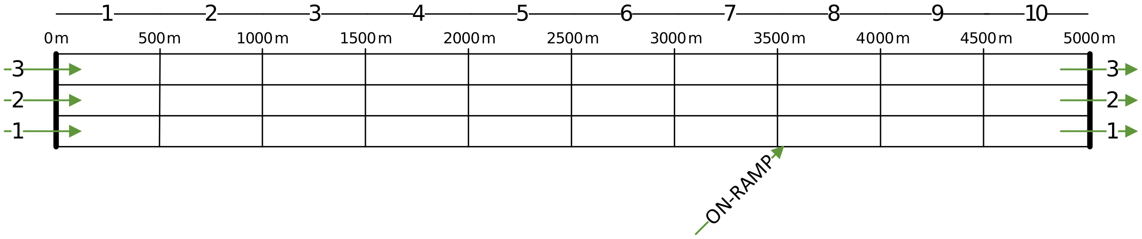

Fig. 5. The motorway stretch used as test-bed for microscopic simulations.

has been tested on a set of 25 replications; different replications share the same mean values for all the stochastic simulation parameters, however a different random seed for each replication generates different respective realisations.

Section 4.2 describes the results obtained in case no control is applied. The next two Sections 4.3 and 4.4 are related to the following respective scenarios: (a) the optimisation layer is applied by itself; and (b) in combination with the local control layer; in both scenarios, it is assumed that all vehicles are equipped and capable of performing the ordered control actions. Section 4.5 reports results obtained with different penetration rates of equipped vehicles, in order to demonstrate and evaluate the robustness of the proposed approach. A final discussion (Section 4.6) summarises and highlights some aspects related to performed simulations. Averaged aggregated numerical results for all the tested scenarios are presented in Table 1.

4.2. Scenario 1: No-control case

Scenario 1 represents the reference case, in which control actions are not applied; it will be used for performance evaluation of the proposed control strategy. A fixed speed limit is set as $v^{max} = 100$ km/h for all the motorway sections and lanes, whereas the lateral movements are delegated to the decisions of the microscopic lane-changing behaviour model, properly modified as described in Section 3.1 to reflect a more realistic merging behaviour.

This scenario is characterised by a strong congestion that starts at the on-ramp merge segment, and covers an upstream part of the simulated motorway stretch. The corresponding TTS (averaged for all the replications) is $TTS = 231.5$ veh h; the standard deviation computed for the set of 25 replications is 14.2 veh h. In the remaining part of this section, a more detailed description of the traffic conditions is presented, related to one single replication with TTS close to the average value.

As it can be seen from the contour plots (by lane) in Fig. 7, congestion is created at lane 1 of the merge segment 8 after about 16 min because of the high demand both from the mainstream and the on-ramp. The congestion quickly spreads over the three lanes and spills back, reaching up to Section 5. At $t = 30$ min, the demand starts to decrease (Fig. 6), leading to the gradual decrease of the congestion extent and its complete disappearance from all lanes after $t = 40$ min.

4.3. Scenario 2: Application of optimal results

According to the topology of the motorway stretch, the macroscopic traffic flow model described in Section 2.3, that is used in the optimisation layer, has been calibrated against the microscopic simulator following the methodology discussed by Roncoli et al. (2015a). This resulted in a reasonable match of the congestion pattern, that comprises also the capacity drop phenomenon.

An optimisation horizon of 10 min has been set for the MPC problem within the optimisation layer, that corresponds to the time necessary to drive along the whole stretch at a speed of 30 km/h, which is a reasonable assumption according to Burger et al. (2013). The update period for the optimisation problem is 1 min. The demand during the optimisation horizon

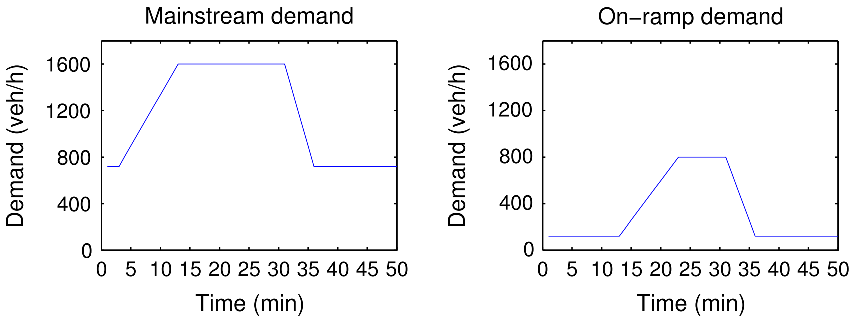

Fig. 6. The trapezoidal demand entering the network origins; for the three lanes of the mainstream, the same mean values are used.

Table 1
Comparison of cases: averages of 25 replications.

	No control	Direct optimal controller				Local control layer			
		100%	30%	10%	5%	100%	30%	10%	5%
Average TTS [veh h]	231.5	210.9	214.6	219.7	219.8	203.4	217.1	219.7	221.2
Improvement (%)	–	9.1	6.1	4.9	3.3	12.1	6.2	5.1	4.4
TTS standard deviation [veh h]	14.2	8.9	13.1	13.1	13.9	7.9	11.3	12	12.5
Improvement (%)	–	37.3	7.3	7.2	2.0	44.5	20.5	15.4	11.5

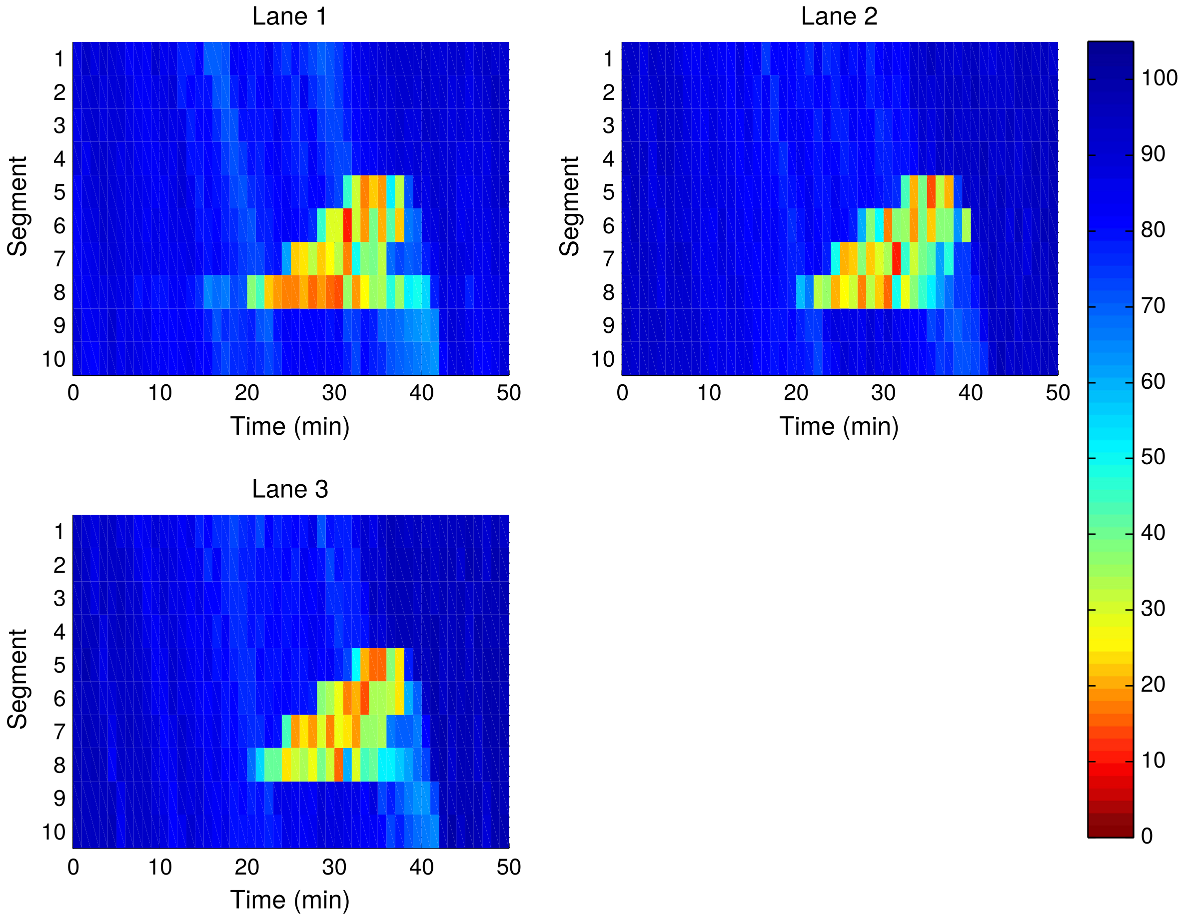

Fig. 7. Contour plots for the mean speed (km/h) in the no-control case (Scenario 1).

is set using the average between historical data (in this case, the mean values shown in Fig. 6) and a constant value corresponding to an exponential smoothing of the actual measured demand.

The control steps for RM and LCC are set to 1 min, whereas for MTFC a shorter time is chosen. The latter choice is due to the following reason: within the optimisation problem, the longitudinal flows are constrained by linear functions in dependence of the current densities; if the control step for MTFC includes more than one model steps, the longitudinal flow values during the entire control step are upper-bounded by all the constraints defined for each one of the comprised model steps; in a dynamic settings, where densities change from model step to model step, this may cause an unreasonable over-constraining of the computed flow values. To avoid this, a different approach is considered: within the optimisation problem, the control step is set equal to the model step (5 s), thus varying synchronously with the problem state variables (i.e. the densities); subsequently, the computed optimal flow values are averaged according to the desired control application step, in this case 1 min, so as to apply all control actions synchronously.

The maximum admissible queue length for RM is 30 veh. This value is considered within the optimisation problem, nevertheless, since the demand cannot be perfectly predicted, peaks of demand may generate an unexpected surge of vehicles queuing. Thus, to avoid a spillover of the ramp queue, a queue management algorithm (Spiliopoulou et al., 2010) is applied at the on-ramp which may override the optimal RM decisions, if necessary, during the simulation. All the weighting coefficients of the optimisation cost function were tuned and kept constant during the entire simulation; in particular, a lower cost for lane-changing is set for segments 6 and 7, from lane 1 towards lane 2, thus encouraging equipped vehicles to anticipate the lane-changes that are expected due to the subsequent merging on-ramp.

This configuration allows to solve the optimisation problem in a computation time between 2 s and 3 s for all the instances of the problem (wall-clock time using an Intel® Core i5 personal computer), which is much lower than the update period (1 min) of the optimisation and hence readily feasible also for real-time applications.

The described methodology, applied to the designed scenario, generates an amelioration of the traffic conditions; specifically, an average $TTS = 210.3$ veh h is obtained, which is a 9.1% improvement with respect to the no-control case; of course, the improvement is higher within a tighter space–time window that includes the congestion. In this Scenario, the standard deviation computed among the 25 replications is 8.9 veh h, which is a 37.3% improvement with respect to the no-control case; this implies that the variation of traffic conditions from replication to replication is relatively strong when no control is applied, but much less pronounced in the control case. Making the rough but reasonable assumption that each replication corresponds to a working day, this result implies a large potential enhancement in terms of travel time reliability, which is a significant objective of modern traffic control systems, as it entails improved predictability of the daily travel times for the road users. Again, the behaviour of one single replication is detailed hereafter.

The main reason beneath the improvement in the control case lies in the mitigation of the congestion-induced capacity drop and the better usage of the three lanes, thanks to the pertinent control actions; this leads to earlier arrivals of vehicles at the network exit, as it can be seen by inspection of Fig. 8; in fact, the overall throughput is seen to be higher in the controlled case during the peak period.

These results are achieved via integrated and synergistic application of all the three available control actions. Specifically:

- A strong RM action is performed during the peak demand period (between $t = 20$ min and $t = 35$ min). Because of the limited storage space on the ramp, after $t = 22$ min, the outflow of the ramp has to be increased; therefore the RM action is not sufficient to fully avoid the congestion. Fig. 9a displays the queue generated at the ramp.
- Appropriate LCC actions take place in segment 6 from lane 1 towards lane 2 (that is characterised by a lower penalty cost for lateral flow) in order to facilitate vehicles entering from the ramp to merge and avoid an excessive increase of vehicles in the merge area; corresponding values are shown in Fig. 9b. Also, more generally, LCC is responsible for maximising the motorway's cross-section capacity via appropriate change-lane orders.
- Because of the full on-ramp, MTFC actions are performed in lanes 1 and 2 (and to a lesser extent also in lane 3) of segments 1–7 from $t = 22$ min to $t = 40$ min, which limit the flow arriving in the merge area. This creates, as it is shown in Fig. 10, a slight controlled congestion, which has a higher internal speed than the one present in the no-control case and, most importantly, a higher throughput thanks to the mitigation of the capacity drop at the merge area. It is worth noting that, since no off-ramps are present in the area upstream of the bottleneck area within the proposed example, the impact of RM and MTFC on the resulting efficiency is basically equivalent, since the related TTS improvement is only due to the avoidance of capacity drop.

4.4. Scenario 3: Introduction of a local control layer

The introduction of a local control layer envisioned in Section 2.4 has been tested and compared with the results previously obtained. The configuration of the optimisation problem is the same described in Section 4.3; however, in this case, the resulting trajectories of density values are used as set-points for local controllers. The feedback controllers are characterised by a control step of 20 s, and consequently the set-points are adjusted according to the corresponding values retrieved from the optimisation, extracting the optimal state trajectories (densities) computed for the corresponding time step, thus changing the set-points every 20 s. For the RM controller, a gain $K_r = 30$ km/h is used; whereas for VSL, the gain is dependent on the speed limit currently applied (gain scheduling), according to experimental findings reported in Muller et al. (2015). Specifically:

$$K_v = \begin{cases} 1.3\ \frac{\text{km}^2}{\text{veh h}}, & \text{if } v_{i,j}(k-1) > 40\ \text{km/h} \\ 0.34\ \frac{\text{km}^2}{\text{veh h}}, & \text{if } 15\ \text{km/h} < v_{i,j}(k-1) \leqslant 40\ \text{km/h} \\ 0.13\ \frac{\text{km}^2}{\text{veh h}}, & \text{if } v_{i,j}(k-1) \leqslant 15\ \text{km/h} \end{cases} \tag{4}$$

The resulting RM and VSL values are subsequently bounded, and the queue management algorithm (Spiliopoulou et al., 2010) is applied at the ramp to ensure that the created on-ramp queue will not exceed its maximum admissible limit.

The resulting traffic pattern can be inspected in Fig. 12; it results in a $TTS = 203.4$ veh h, that is a reduction of 12.1% with respect to the no-control case, and marks also a further improvement with respect to the direct application of optimal results. Also the TTS standard deviation across the 25 replications has improved, obtaining a value of 7.9 veh h, which is a 44.5% improvement with respect to the no-control case, and marks also a further improvement in terms of travel time reliability.

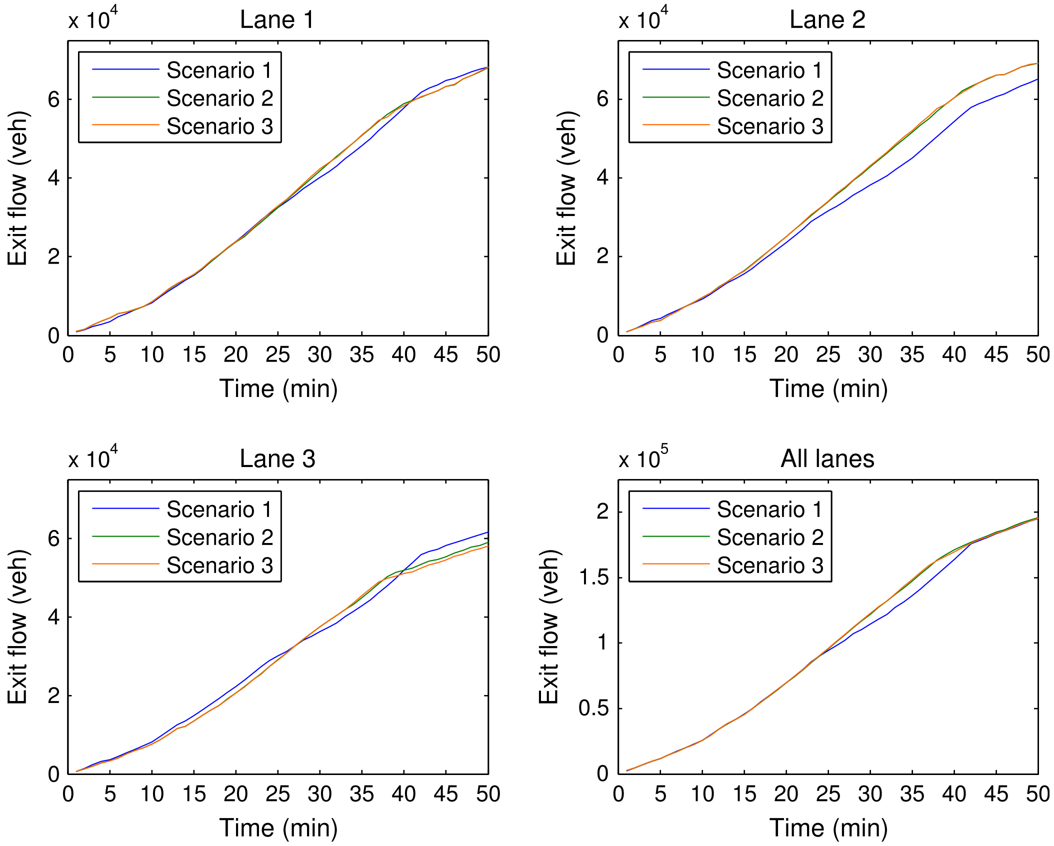

Fig. 8. The time-accumulated flows at the network exit are shown; in the controlled cases, the flow exiting from all lanes during the peak period is increased, which results in an overall TTS improvement. Notice also the slightly changed lane distribution, mainly due to lane-changing control.

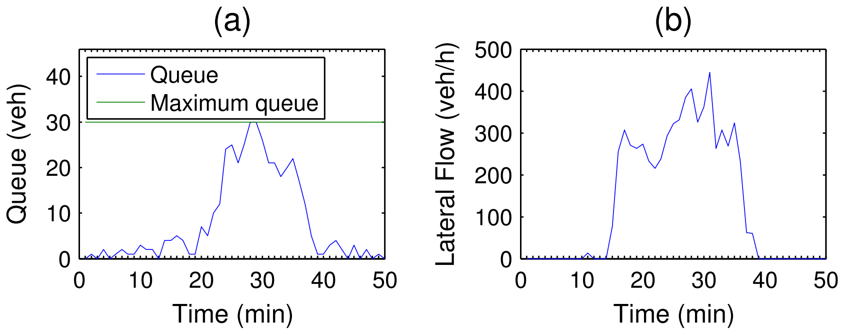

Fig. 9. (a) The queue generated at the on-ramp; and (b) the lateral flow assigned to segment 6, from lane 1 to lane 2, within Scenario 2.

Analysing again one single replication, the control actions appear quite similar to the previous case: the congestion is tackled firstly performing RM actions (Fig. 11a) and then, when the queue is approaching its maximum size (and some queue management actions are taken), also VSL is applied (see contour plots in Fig. 12); LCC is performed, similarly to the previous described case, in segment 6 (see Fig. 11b).

4.5. Mixed traffic

The results presented in the previous sections refer to scenarios where all vehicles are capable to receive and apply the assigned control tasks. In addition, a set of simulations in which different penetration rates are considered (i.e. 30%, 10%, and 5%), is performed in this section. In each defined scenario, only a limited amount of vehicles is capable of receiving and accomplishing the control tasks. Note that this situation does not affect RM, since a traffic light is utilised at the head of

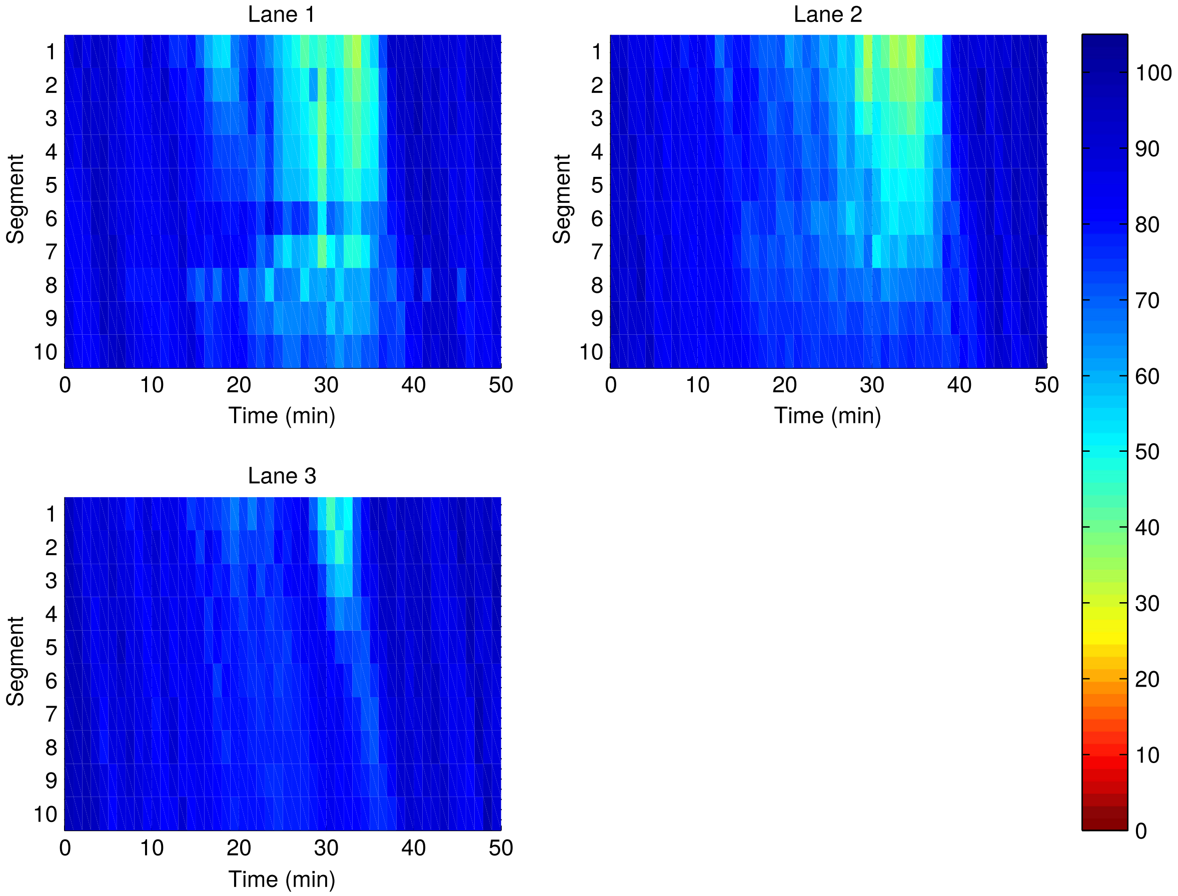

Fig. 10. Contour plots for the mean speed (km/h) in case the results of the optimisation problem are directly applied (Scenario 2).

the on-ramp. For LCC, VACS-equipped vehicles are supposed to change lane only in case the specific order is received, whereas manual vehicles are subject to the lane changing model implemented in the microscopic simulator (i.e. the modified Gipps model described in Section 3.1). Finally, only equipped vehicles apply the VSL received by the control system, while the remaining vehicles move according to the microscopic car-following model and a maximum speed (set as speed limit) of 100 km/h. On the other hand, the traffic state necessary to feed the optimiser and the local controller sis supposed to be exact; this assumption is reasonable in case an appropriate state estimation algorithm is exploited (e.g., Bekiaris-Liberis et al., 2015).

As expected, a reduced penetration rate leads to an increase of the obtained TTS; however, for all the tested scenarios, the TTS is improved with respect to the no-control case. In contrast, standard deviation improvements are not high even for the 30% case, particularly for the direct application of optimal results. Again, corresponding results are included in Table 1.

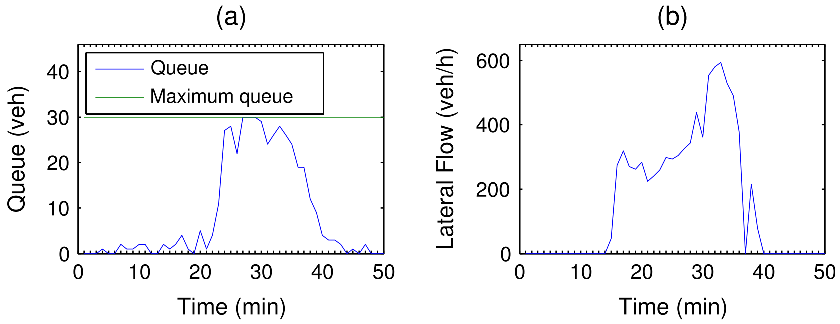

Fig. 11. The queue generated at the on-ramp (a) and the lateral flow assigned to segment 6 (b) within Scenario 3.

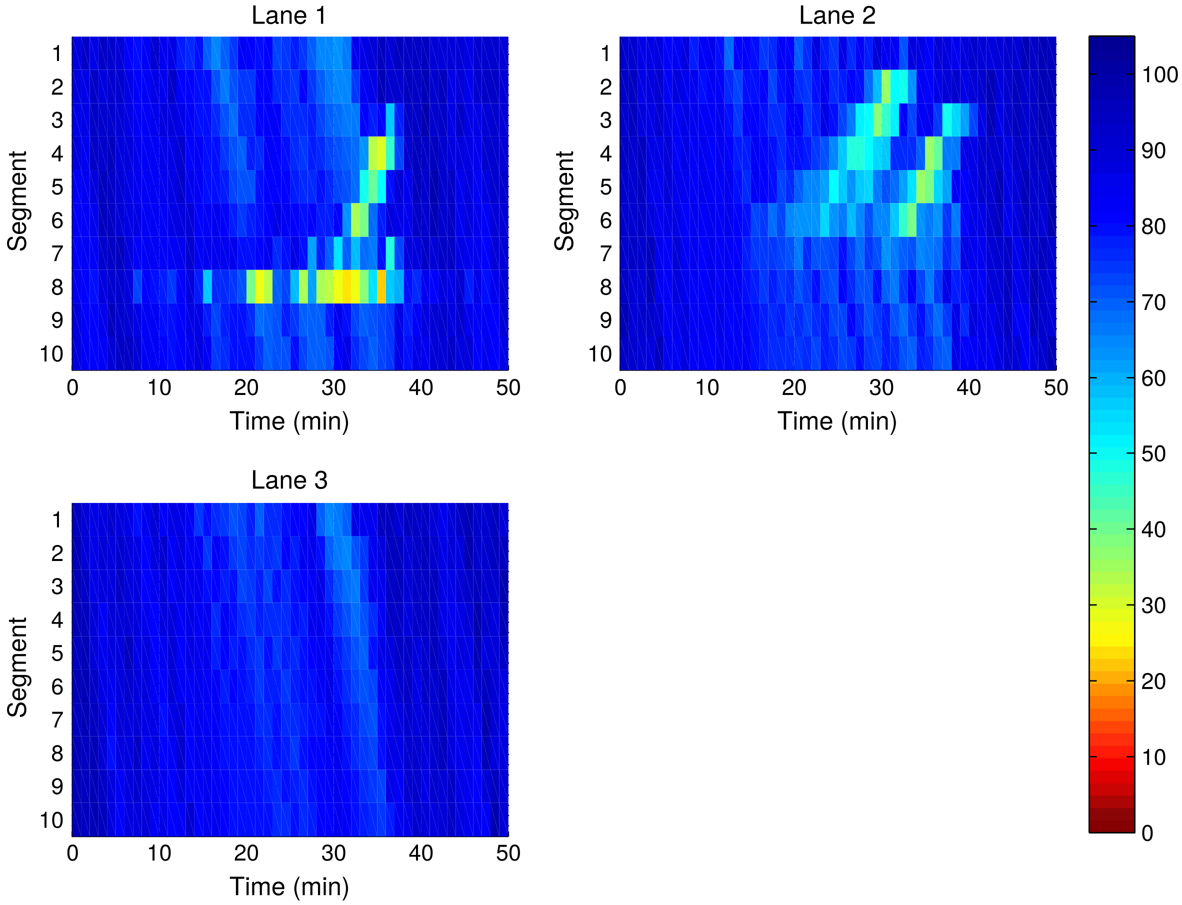

Fig. 12. Contour plots for the mean speed (km/h) in case a local control layer is added (Scenario 3).

The results support the conjecture that, even in the case of a reduced penetration rate of connected vehicles, the proposed methodology can be effective and bring to a consistent amelioration of the traffic conditions. Of course, for decreasing penetration rate, the TTS improvement decreases; however, even for the case with a penetration rate of 5%, the resulting congestion is less severe.

4.6. Discussion

According to the obtained results, the introduced local control layer led to only moderately better control results, e.g. in terms of TTS. This is not surprising in view of the relatively short control time step used for the MPC (1 min) and the relatively small and simple test network; which allow for the MPC to promptly react to external disturbances even in absence of local feedback controllers. More significant additional improvements due to the local control layer are expected for larger networks and more complex traffic conditions. On the other hand, the TTS standard deviation (Table 1) is clearly reduced when a local control layer is present, improving the travel time reliability. Based on these results, we conclude that the introduction of a local control layer, though useful, may not be strictly necessary, as long as the computation complexity of the optimisation problem allows to maintain a sufficiently small update period for the MPC.

As a last remark, in the controlled case, a potential source of TTS degradation may be the application of (minor) control actions in uncongested conditions due to imprecise measurements, model mismatch, or inaccurate numerical approximations. This may be overcome via the definition of an activation/deactivation logic (e.g., using appropriately defined density thresholds), which would permit to apply control only when it is necessary, leaving the system uncontrolled when control actions are not actually needed.

5. Conclusions

The paper presents an MPC approach for solving a coordinated and integrated motorway traffic control problem in presence of VACS-equipped vehicles. The control structure is defined in order to deal with the different aspects of the problem, particularly focusing on the beneficial aspects that the use of VACS could bring to traffic conditions. The chosen convex QP problem facilitates a real-time feasible tool for optimising the proposed coordinated and integrated traffic problem, that can be applied also for large-scale systems. The method calls for very low computation times and guarantees a global optimum, in contrast to other non-linear approaches. The results obtained via microscopic simulation demonstrate that this approach may generate significant improvements in terms of mitigation of traffic congestion, in an application setting where all or a percentage of vehicles were assumed to be equipped with specific devices and to be able to accomplish the given control tasks.

Relevant aspects that are currently under investigation include the consideration of V2I/I2V communication delays for VACS-equipped vehicles, as well as the evaluation of more complex mixed traffic scenarios, where vehicles equipped with different ACC or CACC systems are travelling together with manually driven vehicles.

Acknowledgements

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement n. 321132, project TRAMAN21.

References

Baskar, L.D., De Schutter, B., Hellendoorn, H., 2012. Traffic management for automated highway systems using model-based predictive control. IEEE Trans. Intell. Transp. Syst. 13 (2), 838–847.
Bekiaris-Liberis, N., Roncoli, C., Papageorgiou, M., 2015. Highway traffic state estimation with mixed connected and conventional vehicles using speed measurements. In: Proceedings of the IEEE 18th International Conference on Intelligent Transportation Systems (ITSC), pp. 2806–2811.
Bishop, K.J., Wittenmark, B., 1995. Adaptive Control, second ed. Addison-Wesley.
Bishop, R., 2005. Intelligent Vehicle Technology and Trends. Artech House Publishers.
Breton, P., Hegyi, A., De Schutter, B., Hellendoorn, H., 2002. Shock wave elimination/reduction by optimal coordination of variable speed limits. In: Proceedings. The IEEE 5th International Conference on Intelligent Transportation Systems, pp. 225–230.
Burger, M., van den Berg, M., Hegyi, A., De Schutter, B., Hellendoorn, J., 2013. Considerations for model-based traffic control. Transp. Res. Part C 35, 1–19.
Camacho, E.F., Bordons, C., 1995. Model Predictive Control in the Process Industry. Springer, London.
Carlson, R.C., Papamichail, I., Papageorgiou, M., 2011. Local feedback-based mainstream traffic flow control on motorways using variable speed limits. IEEE Trans. Intell. Transp. Syst. 12 (4), 1261–1276.
Carlson, R.C., Papamichail, I., Papageorgiou, M., 2014. Integrated feedback ramp metering and mainstream traffic flow control on motorways using variable speed limits. Transp. Res. Part C: Emerg. Technol. 46, 209–221.
Carlson, R.C., Papamichail, I., Papageorgiou, M., Messmer, A., 2010. Optimal motorway traffic flow control involving variable speed limits and ramp metering. Transp. Sci. 44 (2), 238–253.
Chen, D., Ahn, S., Hegyi, A., 2014. Variable speed limit control for steady and oscillatory queues at fixed freeway bottlenecks. Transp. Res. Part B: Methodol. 70 (0), 340–358.
Chevallier, E., Leclercq, L., 2009. Do microscopic merging models reproduce the observed priority sharing ratio in congestion? Transp. Res. Part C 17 (3), 328–336.
Chow, A.H., 2015. Optimisation of dynamic motorway traffic via a parsimonious and decentralised approach. Transp. Res. Part C: Emerg. Technol. 55, 69–84.
Daganzo, C.F., 1994. The cell transmission model: a dynamic representation of highway traffic consistent with the hydrodynamic theory. Transp. Res. Part B: Methodol. 28 (4), 269–287.
de Fabritiis, C., Ragona, R., Valenti, G., 2008. Traffic estimation and prediction based on real time floating car data. In: Proceedings of the 11th International IEEE Conference on Intelligent Transportation Systems, pp. 197–203.
Diakaki, C., Papageorgiou, M., Papamichail, I., Nikolos, I.K., 2015. Overview and analysis of vehicle automation and communication systems from a motorway traffic management perspective. Transp. Res. Part A: Policy Pract. 75, 147–165.
Ferrara, A., Sacone, S., Siri, S., 2015. Event-triggered model predictive schemes for freeway traffic control. Transp. Res. Part C: Emerg. Technol. 58, 554–567.
Findeisen, W., Bailey, F.N., Brdys, M., Malinowski, K., Tatjewski, P., Wozniak, A., 1980. Control and Coordination in Hierarchical Systems. John Wiley & Sons.
Frejo, J.R.D., Nez, A., Schutter, B.D., Camacho, E.F., 2014. Hybrid model predictive control for freeway traffic using discrete speed limit signals. Transp. Res. Part C: Emerg. Technol. 46, 309–325.
Gipps, P., 1981. A behavioural car-following model for computer simulation. Transp. Res. Part B 15 (2), 105–111.
Gipps, P., 1986. A model for the structure of lane-changing decisions. Transp. Res. Part B 20 (5), 403–414.
Gomes, G., Horowitz, R., 2006. Optimal freeway ramp metering using the asymmetric cell transmission model. Transp. Res. Part C: Emerg. Technol. 14 (4), 244–262.
Goñi Ros, B., Knoop, V., Van Arem, B., Hoogendoorn, S., 2014. Mainstream traffic flow control at sags. In: Proceedings of the 93th Annual Meeting of the Transportation Research Board (TRB), Washington, D.C., USA.
Hall, R.W., Lotspeich, D., 1996. Optimized lane assignment on an automated highway. Transp. Res. Part C 4 (4), 211–229.
Hegyi, A., De Schutter, B., Hellendoorn, H., 2005a. Model predictive control for optimal coordination of ramp metering and variable speed limits. Transp. Res. Part C: Emerg. Technol. 13 (3), 185–209.
Hegyi, A., De Schutter, B., Hellendoorn, J., 2005b. Optimal coordination of variable speed limits to suppress shock waves. IEEE Trans. Intell. Transp. Syst. 6 (1), 102–112.
Hegyi, A., Hoogendoorn, S., Schreuder, M., Stoelhorst, H., Viti, F., 2008. Specialist: a dynamic speed limit control algorithm based on shock wave theory. In: Proceedings of the 11th International IEEE Conference on Intelligent Transportation Systems, pp. 827–832.
Herrera, J.C., Bayen, A.M., 2010. Incorporation of lagrangian measurements in freeway traffic state estimation. Transp. Res. Part B: Methodol. 44 (4), 460–481.
Kesting, A., Treiber, M., Schönhof, M., Helbing, D., 2007. Extending adaptive cruise control to adaptive driving strategies. Transp. Res. Rec.: J. Transp. Res. Board 2000, 16–24.
Kim, K., Medanić, J.V., Cho, D.I., 2008. Lane assignment problem using a genetic algorithm in the automated highway systems. Int. J. Automot. Technol. 9 (3), 353–364.
Kotsialos, A., Papageorgiou, M., Mangeas, M., Haj-Salem, H., 2002. Coordinated and integrated control of motorway networks via non-linear optimal control. Transp. Res. Part C 10 (1), 65–84.
Lebacque, J., 2003. Two-phase bounded-acceleration traffic flow model: analytical solutions and applications. Transp. Res. Rec. 1852 (1), 220–230.
Li, Y., Chow, A., Cassel, D., 2014. Optimal control of motorways by ramp metering, variable speed limits, and hard-shoulder running. Transp. Res. Rec.: J. Transp. Res. Board 2470, 122–130.

Lighthill, M.J., Whitham, G.B., 1955. On kinematic waves. II. A theory of traffic flow on long crowded roads. Proc. Roy. Soc. A: Math. Phys. Eng. Sci. 229 (1178), 317–345.

Muller, E.R., Carlson, R.C., Kraus, W., Papageorgiou, M., 2015. Microsimulation analysis of practical aspects of traffic control with variable speed limits. IEEE Trans. Intell. Transp. Syst. 16 (1), 512–523.

Muralidharan, A., Horowitz, R., 2012. Optimal control of freeway networks based on the link node cell transmission model. In: Proceedings of the American Control Conference, pp. 5769–5774.

Nissan, A., Koutsopoulos, H.N., 2011. Evaluation of the impact of advisory variable speed limits on motorway capacity and level of service. Procedia – Soc. Behav. Sci. 16, 100–109.

Ntousakis, I.A., Nikolos, I.K., Papageorgiou, M., 2015. On microscopic modelling of adaptive cruise control systems. Transp. Res. Procedia 6, 111–127.

Papageorgiou, M., 1984. Multilayer control system design applied to freeway traffic. IEEE Trans. Autom. Control 29 (6), 482–490.

Papageorgiou, M., Diakaki, C., Dinopoulou, V., Kotsialos, A., Wang, Y., 2003. Review of road traffic control strategies. In: Proceedings of the IEEE, vol. 91, pp. 2043–2065.

Papageorgiou, M., Hadj-Salem, H., Blosseville, J.-M., 1991. ALINEA: a local feedback control law for on-ramp metering. Transp. Res. Rec. 1320, 58–64.

Papageorgiou, M., Papamichail, I., 2008. Overview of traffic signal operation policies for ramp metering. Transp. Res. Rec. 2047, 28–36.

Papamichail, I., Kotsialos, A., Margonis, I., Papageorgiou, M., 2010. Coordinated ramp metering for freeway networks a model-predictive hierarchical control approach. Transp. Res. Part C 18 (3), 311–331.

Papamichail, I., Papageorgiou, M., Vong, V., Gaffney, J., 2010. Heuristic ramp-metering coordination strategy implemented at Monash Freeway, Australia. Transp. Res. Rec. 2178, 10–20.

Ramaswamy, D., Medanic, J., Perkins, W., Benekohal, R., 1997. Lane assignment on automated highway systems. IEEE Trans. Veh. Technol. 46 (3), 755–769.

Rao, B., Varaiya, P., 1994. Roadside intelligence for flow control in an intelligent vehicle and highway system. Transp. Res. Part C 2 (1), 49–72.

Richards, P.I., 1956. Shock waves on the highway. Oper. Res. 4 (1), 42–51.

Roncoli, C., Papageorgiou, M., Papamichail, I., 2015a. Traffic flow optimisation in presence of vehicle automation and communication systems – Part I: a first-order multi-lane model for motorway traffic. Transp. Res. Part C 57, 241–259.

Roncoli, C., Papageorgiou, M., Papamichail, I., 2015b. Traffic flow optimisation in presence of vehicle automation and communication systems – Part II: optimal control for multi-lane motorways. Transp. Res. Part C 57, 260–275.

Roncoli, C., Papamichail, I., Papageorgiou, M., 2014. Model predictive control for multi-lane motorways in presence of VACS. In: Proceedings of the IEEE 17th International Conference on Intelligent Transportation Systems, pp. 501–507.

Seo, T., Kusakabe, T., Asakura, Y., 2014. Estimation of flow and density using probe vehicles with spacing measurement equipment. Transp. Res. Part C 53, 134–150.

Smaragdis, E., Papageorgiou, M., 2003. Series of new local ramp metering strategies. Transp. Res. Rec. 1856, 74–86.

Smulders, S., Helleman, D., 1998. Variable speed control: state-of-the-art and synthesis. In: Proceedings of the 9th International Conference on Road transport information and control, pp. 155–159.

Spiliopoulou, A.D., Manolis, D., Papamichail, I., Papageorgiou, M., 2010. Queue management techniques for metered freeway on-ramps. Transp. Res. Rec. 2178, 40–48.

Sun, X., Horowitz, R., 2005. A localized switching ramp-metering controller with a queue length regulator for congested freeways. In: Proceedings of the 2005 American Control Conference, vol. 3, pp. 2141–2146.

Torné, J., Soriguera, F., Geroliminis, N., 2014. Coordinated active traffic management freeway strategies using capacity-lagged cell transmission model. In: Proceedings of the TRB 93rd Annual Meeting, Washington, D.C., USA.

Treiber, M., Hennecke, A., Helbing, D., 2000. Congested traffic states in empirical observations and microscopic simulations. Phys. Rev. E 62 (2), 1805–1824.

TSS – Transport Simulation Systems, January 2014. Aimsun 8 Users Manual.

Van Lint, J., Hoogendoorn, S.P., 2010. A robust and efficient method for fusing heterogeneous data from traffic sensors on freeways. Comput.-Aided Civ. Infrastruct. Eng. 25 (8), 596–612.

VanderWerf, J., Shladover, S., Kourjanskaia, N., Miller, M., Krishnan, H., 2001. Modeling effects of driver control assistance systems on traffic. Transp. Res. Rec. 1748 (01), 167–174.

Varaiya, P., 1993. Smart cars on smart roads: problems of control. IEEE Trans. Autom. Control 38 (2), 195–207.

Wang, J., Liu, R., Montgomery, F., 2005. Car-following model for motorway traffic. Transp. Res. Rec. 1934, 33–42.

Yang, X., Lin, Y., Lu, Y., Zou, N., 2013. Optimal variable speed limit control for real-time freeway congestions. Procedia – Soc. Behav. Sci. 96, 2362–2372.

Yuan, Y., van Lint, J., Wilson, R., van Wageningen-Kessels, F., Hoogendoorn, S., 2012. Real-time lagrangian traffic state estimator for freeways. IEEE Trans. Intell. Transp. Syst. 13 (1), 59–70.

Zhang, J., Chang, H., Ioannou, P., 2006. A simple roadway control system for freeway traffic. In: Proceedings of the 2006 American Control Conference, pp. 4900–4905.

Zhou, X., Mahmassani, H.S., 2007. A structural state space model for real-time traffic origin–destination demand estimation and prediction in a day-to-day learning framework. Transp. Res. Part B: Methodol. 41 (8), 823–840.