

Technical report CSE02-023

Optimal coordination of variable speed limits to suppress shock waves*

A. Hegyi, B. De Schutter, and J. Hellendoorn

If you want to cite this report, please use the following reference instead:

A. Hegyi, B. De Schutter, and J. Hellendoorn, "Optimal coordination of variable speed limits to suppress shock waves," *IEEE Transactions on Intelligent Transportation Systems*, vol. 6, no. 1, pp. 102–112, Mar. 2005.

Control Systems Engineering
Faculty of Information Technology and Systems
Delft University of Technology
Delft, The Netherlands
Current URL: <https://www.dcsc.tudelft.nl>

*This report can also be downloaded via https://pub.deschutter.info/abs/02_023.html

Optimal Coordination of Variable Speed Limits to Suppress Shock Waves

A. Hegyi, B. De Schutter and J. Hellendoorn

Abstract—When freeway traffic is dense, shock waves may appear. These shock waves result in longer travel times and in sudden, large variations in the speeds of the vehicles, which could lead to unsafe situations. Dynamic speed limits can be used to eliminate or at least to reduce the effects of shock waves. However, coordination of the variable speed limits is necessary in order to prevent the occurrence of new shock waves and/or a negative impact on the traffic flows in other locations. In this paper we present a model predictive control (MPC) approach to optimally coordinate variable speed limits for freeway traffic with the aim of suppressing shock waves. First, we optimize continuous valued speed limits, such that the total travel time is minimal. Next, we include a safety constraint that prevents drivers from encountering speed limit drops larger than, e.g., 10 km/h. Furthermore, to get a better correspondence between the computed and the applied control signals, we now consider discrete speed limits. We illustrate our approach with a benchmark problem.

Index Terms—traffic control, coordinated control, variable speed limits, safe speed limits, model predictive control, shock waves.

I. INTRODUCTION

A. Problem description

AS the number of vehicles grows and the need for mobility increases on a world-wide scale, the frequency and duration of traffic jams in and around major cities increase. In the short term the most effective measures in the battle against traffic congestion seem to be a selective construction of new roads — an option which is often not viable due to lack of space and/or budgetary means, — and a more efficient use of the existing infrastructure through dynamic traffic management and control. One of the strategies used in dynamic traffic

A. Hegyi, B. De Schutter and J. Hellendoorn are with Delft Center for Systems and Control, Delft University of Technology, Mekelweg 2, 2628 CD, Delft, The Netherlands, phone: +31-15-278.20.87, fax: +31-15-278.66.79, {A.Hegyi,B.DeSchutter,J.Hellendoorn}@DCSC.TUdelft.nl
This research was supported by the Traffic Research Centre (AVV) of the Dutch Ministry of Transport, Public Works and Water Management and by the *Mobility of People and Transportation of Goods* spearhead program of the Delft University of Technology.

Manuscript received:

revised:

control is traffic signal control to obtain a smoother or better circulation of traffic in urban areas, or on freeways traffic conditions are improved by regulating access to freeways or main roads (ramp metering), influencing route choice (variable message signs), and reducing (the probability of) congestion by the use of speed limits. In this paper we will concentrate on dynamic traffic control to fight congestion.

In practice, ramp metering set-ups and other dynamic traffic management installations usually operate based on local data only. However, considering the effects of the measures on the network level instead, and computing the control signals based on network-wide measurements and predictions offers many advantages. Therefore, we will consider a network-wide *coordination* of control measures based on network-wide data. Since the effect of a control measure on more distant locations might only be visible after some time, a *prediction* of the network evolution is also necessary to achieve optimal network control. The model predictive control approach to dynamic traffic management presented in this paper contains both elements: network-wide coordination and prediction.

In this paper we consider a special case of traffic control measures: variable speed limits to reduce or eliminate shock waves. Also in this case prediction and coordination are necessary for an effective control strategy. Prediction is needed for two reasons: first, if the formation or the arrival of a shock wave in the controlled area can be predicted, then preventive measures can be taken. Second, the positive effect of speed limits on the traffic flow cannot be observed instantaneously,¹ so prediction is necessary and should at least include the point when the improvement can be observed.

Another advantage of a predictive controller is that it may prevent a certain type of instability. In control systems theory it is well-known that the delay between the appearance of a disturbance and the action of the controller may cause instabilities. By using predictive control the disturbance can be anticipated on, and this

¹We will see that the speed limits have to slow down a part of the traffic first, in order to dissolve the shock wave.

kind of instabilities may be prevented.

Besides prediction and coordination, the speed limit control problem has other characteristics that impose certain requirements on the control strategy:

- 1) There is a direct relation between the outflow of a network and the total time spent (TTS) in the network, assuming that the traffic demand is fixed. Papageorgiou *et al.* [1] showed that in a traffic network an increase of outflow of 5% may result in an decrease of the total time spent in the network of 20%². Since the congestion after a breakdown usually has an outflow that is 5–10% lower than the capacity³ we can conclude that any control method that resolves (reduces) congestion will at best achieve a flow improvement of approximately 5–10%, but this improvement can decrease the TTS significantly. This also means that the control strategy requires a high precision. For this reason and because there are always (unpredictable) disturbances present in a traffic network, *feedback control* is required.
- 2) The speed limit signs used in practice display speed limits in increments of e.g. 10 or 20 km/h. Therefore, the controller should produce discrete control signals.⁴
- 3) For safety it is often required that the drivers should not encounter a decrease in the displayed speed limit larger than a prespecified amount. The controller should be able to take this kind of constraints into account.

The control strategy presented in this paper takes these three requirements into account.

B. Literature survey

In the literature, basically two views on the use of speed limits can be found. The first emphasizes the homogenization effect [4], [5], [6], [7], [8], whereas the second is more focused on the preventing traffic breakdown by reducing the flow by means of speed limits [9], [10], [11].

The basic idea of homogenization is that speed limits can reduce the speed differences, by which a stabler (and safer) flow can be achieved. The homogenizing

²This result is specific to this example, in other situations the decrease of TTS could be different. Nevertheless, the example clarifies that there is a strong relation between outflow and TTS in case of congested networks.

³This is the so-called capacity-drop phenomenon (see, e.g., [2]).

⁴In this paper the main focus is on the procedure how discrete speed limits are found without solving a computationally demanding discrete optimization problem. A comparative study of different speed limit step sizes can be found in [3].

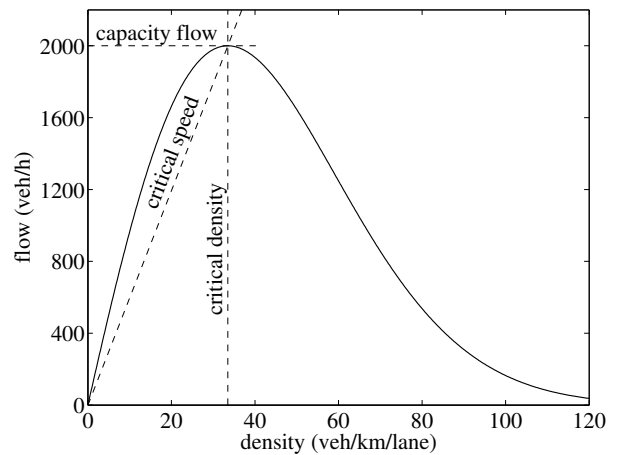


Figure 1. A typical example of the fundamental diagram. The critical speed is the speed that corresponds to the maximum flow. The slope of the line connecting the origin and any point in the diagram represents the speed corresponding to that point.

approach typically uses speed limits that are above the critical speed (i.e., the speed that corresponds to the maximal flow; see Figure 1 for an illustration). So, these speed limits do not limit the traffic flow, but only slightly reduce the average speed (and slightly increase the density). The homogenization approach can in theory increase the time to breakdown [12], but cannot suppress or resolve existing shock waves. According to field tests homogenization results in a somewhat more stable and safer traffic flow, but no significant improvement of traffic volume measured [6], [8] (nor can be expected based on theory). Similar results are reported in the recently completed field test at the freeway A13 near Rotterdam in The Netherlands. The main effect of the homogenizing⁵ speed limits (in this case fixed 80 km/h) is the improvement of air quality and safety [13], [14]. An interesting feature of another field test performed on the freeway A1 near Deventer in The Netherlands [8] is that the homogenizing speed limit was dynamically chosen according to a switching logic as a function of the prevailing traffic conditions. Also this test resulted in similar findings regarding the effects on speed and flow. An extended overview of existing speed limit systems that aim at reducing speed differentials is given by Wilkie [15].

The traffic breakdown prevention approach focuses more on preventing too high densities, and also allows speed limits that are lower than the critical speed in order to limit the inflow to the jammed area. By resolving these high density areas (bottlenecks) higher flow can be achieved in contrast to the homogenization approach.

⁵In this experiment the main goal was to reduce noise and air pollution. However, the applied speed limits resulted in a more homogeneous traffic flow.

Several control methodologies are used in literature to find a control law for speed control, such as multi-layer control [16], sliding-mode control [10], [11], and optimal control [17]. In [18] optimal control is approximated by a neural network in a rolling horizon framework. Other authors use (or simplify their control law to) a control logic where the switching between the speed limit values is based on traffic volume, speed or density measurement [4], [5], [6], [7], [11], [12]. In some cases the switching between the speed limit values is also based on special circumstances, such as weather and light conditions [4].

Some authors recognize the importance of anticipation in the speed control scheme. A pseudo-anticipative scheme is used in [11] by switching between speed limits based on the density of the neighboring downstream segment. Real predictions are used in [17], [18] and this is the only approach that results in a significant flow improvement.

Most application oriented studies [5], [6], [15], [8], [13], [14] enforce speed limits, except for [4]. In [4] Zackor mentions that many drivers considered the speed limit as a recommendation, which has led to a slight increase of the mean speed. Enforcement is usually accepted by the drivers if the speed limit system leads to a more stable traffic flow.

Some approaches do explicitly model the effects of speed limits, but many of the models used in literature represent the speed limits by a factor that down-scales the fundamental diagram, or they assume that speed limits influence traffic states that have a lower speeds than the imposed speed limit (see, e.g., [10], [17]). This can give too optimistic results (see Section III-B2). Therefore, we will propose a new macroscopic speed limit model. We also introduce a model to express the difference in the drivers' anticipation to increasing or decreasing downstream densities in order to better reproduce shock waves.

C. Overview of the paper

This paper extends the results of our conference papers [19], [20], in which we have already demonstrated the effectiveness of model predictive control with continuous speed limits against shock waves. In this paper we include a safety constraint that prevents drivers from encountering speed limit drops larger than, e.g., 10 km/h. Furthermore, to get a better correspondence between the computed and the applied control signals, we now consider discrete speed limits (see also [20]).

The organization of the paper is as follows. In Section II we describe the problem of suppressing or reducing the effects of moving jams, and the basic idea behind

our approach to solve this problem. In Section III the basic ingredients of model predictive control for traffic flow control are introduced, and the prediction model including the extensions is presented. The proposed control method is applied to a benchmark problem in Section IV. Finally, the conclusions and topics for future research are stated in Section V.

II. PROBLEM STATEMENT

It is well known (see, e.g., [21]) that some type of traffic jams move upstream with approximately 15 km/h. These jams can remain stationary for a long time, so every vehicle that enters the freeway upstream of the jammed area will have to pass through the jammed area, which increases the travel time. Besides the increased travel time another disadvantage of the moving jams is that they are potentially unsafe.

To suppress shock waves one can use speed limits in the following way. In some sections upstream of a shock wave speed limits are imposed in order to reduce the inflow to the jammed area. When the inflow of the jammed area is reduced sufficiently, i.e., to a lower value than its outflow, the jam will eventually dissolve. In other words, the speed limits create a low-density wave (with a density lower than that is than in the uncontrolled situation) that propagates downstream. This low-density wave meets and compensates the high density shock wave. As a result, the shock wave is reduced or eliminated.

The general idea how dynamic speed limits can dissolve shock waves can also be explained in terms of the stable, metastable, and unstable traffic flow states observed by Kerner and Rheborn [21]. Stable means that the traffic demand is such that any disturbance (no matter how large) will vanish without intervention. The metastable state correspond with traffic demands where both free-flow traffic and a shock wave can remain existent for a long time. In this state, in free flow, small disturbances will typically vanish, but large disturbances will create a shock wave. Unstable means that any disturbance (no matter how small) will trigger a shock wave. If speed limits are to dissolve the shock waves, the traffic flow must be in the metastable state, because in the stable state the disturbances will appear without control, and in the unstable state any speed limit change will initiate a new shock wave. In the metastable state, the speed limits have the possibility (if the increment of the speed limit values is sufficiently small) to spread out the shock wave into a disturbance that is small enough to vanish automatically.

Note that if the speed limits are optimized properly, they will never create a shock wave or an upstream

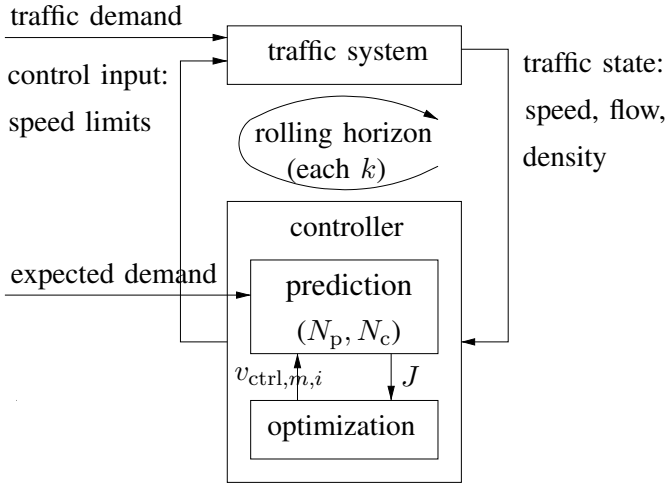


Figure 2. Schematic view of the model predictive control (MPC) structure.

queue that gives rise to delays that are higher than in the uncontrolled case. In this sense “proper” optimization also means that the control objective includes all effects of the speed limits: the effects on the controlled freeway stretch, and the (possible) upstream queues caused by the speed limits.

In the following sections we demonstrate how the proper speed limits can be found.

III. APPROACH

A. Model Predictive Control

We use a model predictive control (MPC) scheme to solve the problem of optimal coordination of speed limits, see Figure 2. In MPC, at each time step k the optimal control signal is computed (by numerical optimization) over a prediction horizon N_p . A control horizon N_c ($< N_p$) is selected to reduce the number of variables and to improve the stability of the system. (After the control horizon has been passed, the control signal is usually taken to be constant.) From the resulting optimal control signal only the first sample $k + 1$ is applied to the process. In the next time step ($k + 1$), a new optimization is performed (with a prediction horizon that is shifted one time step ahead) and of the resulting control signal again only the first sample is applied, and so on. This scheme, called rolling horizon, allows for updating the state from measurements, or even for updating the model in every iteration step.

The advantage of updating the state is that this results in a controller that has a low sensitivity to prediction errors. Regularly updating the model results in an adaptive control system, which could be useful in situations where the model significantly changes, such as in case of incidents or changing weather conditions.

Besides the low sensitivity for prediction errors, this control scheme has another advantage that is a consequence of the prediction that is used. Prediction makes it possible to *perform temporarily sub-optimally in order to gain more performance in the future*. In our case, we will see (what intuitively already can be expected) that in order to suppress a shock wave, traffic flow has to be limited, but when the shock wave has dissolved, the traffic flow will be higher than otherwise.

For more information on MPC see [22], [23], [24] and the references therein.

B. Prediction model

The MPC procedure includes a prediction of the network evolution as a function of the current state and a given control input. For this prediction we use an extended version of the (destination-independent version) of the macroscopic traffic flow model METANET [25], [26], [27]. The extensions are introduced to model shock waves better and to include the effects of speed limits.

Note that the MPC approach is generic, also other traffic models (that include the effect of the speed limit) could be used. The METANET model has been validated in several studies [28], [29], [30]. The reported validation results are in general satisfactory, except for the results in [29], which can be explained by the fact that the model in [29] was not calibrated before validation.

For the sake of brevity, we describe only those parts of the model that are relevant for interpreting and understanding the simulation results of our benchmark network (see Section IV). See [25], [26], [27] for the complete METANET model (incl. lane drops, merging and weaving processes).

1) *Original METANET model*: The METANET model represents a network as a directed graph with the links corresponding to freeway stretches. Each freeway link has uniform characteristics, i.e., no on-ramps or off-ramps and no major changes in geometry. Each link m is divided into N_m segments of length L_m and has λ_m lanes (see Figure 3). Each segment i of link m is characterized by the *traffic density* $\rho_{m,i}(k)$ (veh/lane/km), the *mean speed* $v_{m,i}(k)$ (km/h), and the *traffic volume or flow* $q_{m,i}(k)$ (veh/h), where k indicates the time instant $t = kT$, and T is the time step used for the simulation of the traffic flow (typically $T = 10$ s). The following equations describe the evolution of the network over time. The fundamental relationship between speed, flow and density is reads

$$q_{m,i}(k) = \rho_{m,i}(k) v_{m,i}(k) \lambda_m . \quad (1)$$

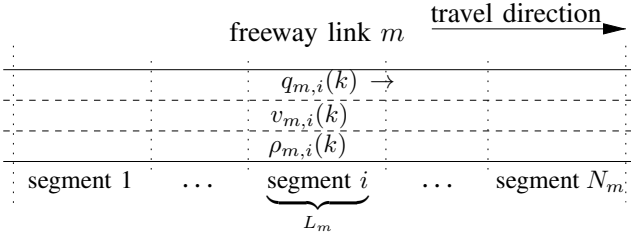


Figure 3. In the METANET model, a freeway link is divided into segments. The main variables in the model are the average outflow of a segment $q_{m,i}(k)$, the average speed $v_{m,i}(k)$, the average density $\rho_{m,i}(k)$, and the segment length L_m .

The conservation of vehicles is expressed by

$$\rho_{m,i}(k+1) = \rho_{m,i}(k) + \frac{T}{L_m \lambda_m} (q_{m,i-1}(k) - q_{m,i}(k)) .$$

The mean speed equals the sum of mean speed at the previous time instance, a relaxation term, a convection term, and an anticipation term that expresses the speed decrease (increase) as drivers experience a density increase (decrease) downstream:

$$\begin{aligned} v_{m,i}(k+1) = & v_{m,i}(k) + \frac{T}{\tau} \left(V(\rho_{m,i}(k)) - v_{m,i}(k) \right) + \\ & \frac{T}{L_m} v_{m,i}(k) (v_{m,i-1}(k) - v_{m,i}(k)) - \\ & \frac{\eta T}{\tau L_m} \frac{\rho_{m,i+1}(k) - \rho_{m,i}(k)}{\rho_{m,i}(k) + \kappa} , \end{aligned} \quad (2)$$

where τ , η and κ are model parameters, and with

$$V(\rho_{m,i}(k)) = v_{\text{free},m} \exp \left[-\frac{1}{a_m} \left(\frac{\rho_{m,i}(k)}{\rho_{\text{crit},m}} \right)^{a_m} \right] , \quad (3)$$

with a_m a model parameter, and where the free-flow speed $v_{\text{free},m}$ is the average speed that drivers would assume if traffic would be flowing freely, and the critical density $\rho_{\text{crit},m}$ is the density at which the traffic flow becomes unstable.

Origins are modeled with a simple queue model. The length of the queue $w_o(k)$ equals the previous queue length plus the demand $d_o(k)$, minus the outflow $q_o(k)$:

$$w_o(k+1) = w_o(k) + T(d_o(k) - q_o(k)) .$$

The outflow depends on the traffic conditions on the freeway and the capacity of the origin. The flow $q_o(k)$ is the minimum of the demand and the maximal flow that can enter the freeway given the mainstream conditions:

$$q_o(k) = \min \left[d_o(k) + \frac{w_o(k)}{T}, Q_o \frac{\rho_{\text{max}} - \rho_{\mu,1}(k)}{\rho_{\text{max}} - \rho_{\text{crit},\mu}} \right] , \quad (4)$$

where Q_o is the on-ramp capacity (veh/h) under free-flow conditions, ρ_{max} is the maximum density, and μ the index of the link to which the on-ramp is connected.

2) *Extensions*: Since the original METANET model does not explicitly describe the effect of speed limits, we have modified the equation for the desired speed (3) to incorporate speed limits. The second extension regards the modeling of a mainstream origin, which has a different nature than an on-ramp origin. The third extension describes the different effects of a positive or negative downstream density gradient on the speed (cf. the anticipation term in (2)).

In some publications, the effect of the speed limit is expressed by downscaling the desired speed-density diagram or using a model in which the speed limits influence traffic states that have a lower speeds than the imposed speed limit (see [10], [17]). This changes the whole speed-density diagram, also for the states where the speed would otherwise be lower than the value of the speed limit. This means, e.g., that if the free flow speed is 120 km/h and the displayed speed limit is 100 km/h then the speed and flow of the traffic are reduced with the factor 100/120 even when the vehicles are traveling at 80 (< 100) km/h. An example of such a speed limit model is the one introduced by Cremer [31], shown in Figure 4 (top), where the fundamental diagram is approximately scaled linearly by the speed limit. Furthermore, scaling down the desired speed also reduces always the capacity, while there is no reason to assume that a speed limit *above* the critical speed (speeds where the flow has not reached capacity yet) would reduce the capacity of the road (see Figure 4 (top)).

The above assumptions are rather unrealistic, and they exaggerate the effect of speed limits. To get a more realistic model for the effects of the speed limits, we assume that the desired speed is the minimum of the following two quantities: the desired speed based on the experienced traffic conditions (in the case of METANET: based on the experienced density), and the desired speed caused by the speed limit displayed on the variable message sign⁶. A similar assumption has been made in [32]. The desired speed equation reads now:

$$\begin{aligned} V(\rho_{m,i}(k)) = & \min \left((1 + \alpha) v_{\text{ctrl},m,i}(k), \right. \\ & \left. v_{\text{free},m} \exp \left[-\frac{1}{a_m} \left(\frac{\rho_{m,i}(k)}{\rho_{\text{crit},m}} \right)^{a_m} \right] \right) , \end{aligned} \quad (5)$$

⁶Although the quantitative validation would give a stronger confirmation of the correctness of this speed limit model, the quantitative arguments presented here and in [3] imply that the proposed model is more adequate than other existing ones. Model validation is a topic for future research.

where $v_{\text{ctrl},m,i}(k)$ is the speed limit imposed on segment i , link m , at time k , and the factor $(1 + \alpha)$ expresses the non-compliance, i.e., the factor that the desired speed is higher than the displayed speed limit. See Figure 4 (bottom) for an illustration. We refer the interested reader to [3] for a more extensive discussion of existing macroscopic speed limit models.

The second extension is introduced to express the different natures of a mainstream origin link o and a regular on-ramp (the queue at a mainstream origin is in fact an abstraction of the sections upstream of the origin of the part of the freeway network that we are modeling). To this end, we use a modified version of (4) with another flow constraint to model a mainstream origin link, because the inflow of a segment (and thus the outflow of the mainstream origin) can be limited by an active speed limit or by the actual speed in the first segment (when either of them is lower than the speed at critical density). Hence, we assume that the maximal flow equals the flow that follows from the speed-flow relationship from (1) and (3) with the speed equal to the speed limit or the actual speed in the first segment, whichever is smaller. So if o is the origin of mainstream link μ , then we have

$$q_o(k) = \min \left[d_o(k) + \frac{w_o(k)}{T}, q_{\text{lim},\mu,1}(k) \right],$$

where $q_{\text{lim},\mu,1}(k)$ is the maximal inflow determined by the limiting speed in the first segment of link μ :

$$q_{\text{lim},\mu,1}(k) = \begin{cases} \lambda_\mu v_{\text{lim},\mu,1}(k) \rho_{\text{crit},\mu} \left[-a_\mu \ln \left(\frac{v_{\text{lim},\mu,1}(k)}{v_{\text{free},m}} \right) \right]^{\frac{1}{a_\mu}} & \text{if } v_{\text{lim},\mu,1}(k) < V(\rho_{\text{crit},\mu}) \\ q_{\text{cap},\mu} & \text{if } v_{\text{lim},\mu,1}(k) \geq V(\rho_{\text{crit},\mu}), \end{cases}$$

where $v_{\text{lim},\mu,1}(k) = \min(v_{\text{ctrl},\mu,1}(k), v_{\mu,1}(k))$ is the speed that limits the flow, and $q_{\text{cap},\mu} = \lambda_\mu V(\rho_{\text{crit},\mu}) \rho_{\text{crit},\mu}$ is the capacity flow.

The third extension was introduced to be able to express the different reactions of drivers to varying downstream densities, since the effect of a higher downstream density is usually stronger than the effect of a lower downstream density. The sensitivity of the speed to the downstream density is expressed by parameter η . In (2) η is a global parameter and has the same value for all segments. However, here we take different values for $\eta_{m,i}(k)$, depending on whether the downstream density is higher or lower than the density in the actual segment:

$$\eta_{m,i}(k) = \begin{cases} \eta_{\text{high}} & \text{if } \rho_{m,i+1}(k) \geq \rho_{m,i}(k) \\ \eta_{\text{low}} & \text{if } \rho_{m,i+1}(k) < \rho_{m,i}(k). \end{cases}$$

C. Objective function

The objective function consist of two terms, a term for the TTS and a term that penalizes large control signal variations. In the TTS term the total number of vehicle-hours spent on the freeway segment and in the origin queue are summed. The control signal variation penalty term is included to express the preference for smooth signals. The trade-off between these two terms is expressed by the relative weighting of these terms in the objective function. Since the primary aim of the controller is to minimize the TTS, the weights are chosen such that the TTS term weighted to be more important.

Note that we distinguish between the controller time step length T_c and the simulation time step length T , and between the controller time step counter k_c and the model time step counter k . We assume that the controller time step length is an integer multiple of the simulation time step length: $T_c = MT$, with M a positive integer.

$$J(k_c) = T \sum_{k=Mk_c}^{M(k_c+N_p)-1} \left\{ \sum_{(m,i) \in I_{\text{links}}} \rho_{m,i}(k) L_m \lambda_m + \sum_{o \in I_{\text{orig}}} w_o(k) \right\} + \alpha_{\text{speed}} \sum_{\ell=k_c}^{k_c+N_c-1} \sum_{(m,i) \in I_{\text{ctrl}}} \left(\frac{v_{\text{ctrl},m,i}(\ell) - v_{\text{ctrl},m,i}(\ell-1)}{v_{\text{free},m}} \right)^2,$$

where I_{links} is the set of indexes of all pairs of segments and links and I_{orig} is the set of all origins, and I_{ctrl} is the set of pairs of indexes (m, i) of the links and segments where speed control is applied. This objective function contains a term for the TTS, and a term that penalizes abrupt variations in the speed limit control signal. The variation term is weighted by the nonnegative weight parameter α_{speed} .

In conventional MPC heuristic tuning rules have been developed to select appropriate values for N_p and N_c [23]. However, these rules cannot be straightforwardly applied to the traffic flow control framework presented in this section. We will select appropriate values for N_p and N_c based on considerations discussed in [33].

D. Constraints

In general, for the safe operation of a speed control system, it is required that the maximum decrease in speed limits that a driver can encounter (v_{maxdiff}) is limited. There are three situations where a driver can encounter a different speed limit value:

- 1) when the speed limit changes in a given segment (and there are more speed limit signs on the same segment),

- 2) when a driver enters a new segment,
- 3) when the driver enters a new segment and the speed limit changes.

The maximum speed difference constraints for the three situations are formulated as follows:

$$\begin{aligned}
v_{\text{ctrl},m,i}(l-1) - v_{\text{ctrl},m,i}(l) &\leq v_{\text{maxdiff}} \\
\text{for all } m, i, l \text{ such that } (m, i) &\in I_{\text{ctrl}} \text{ and} \\
l &\in [k, \dots, k + N_c - 1], \\
v_{\text{ctrl},m,i}(l) - v_{\text{ctrl},m,i+1}(l) &\leq v_{\text{maxdiff}} \\
\text{for all } m, i, l \text{ such that } (m, i) &\in I_{\text{ctrl}}, \\
(m, i+1) &\in I_{\text{ctrl}} \text{ and } l \in [k, \dots, k + N_c - 1], \\
v_{\text{ctrl},m,i}(l-1) - v_{\text{ctrl},m,i+1}(l) &\leq v_{\text{maxdiff}} \\
\text{for all } m, i, l \text{ such that } (m, i) &\in I_{\text{ctrl}}, \\
(m, i+1) &\in I_{\text{ctrl}} \text{ and } l \in [k, \dots, k + N_c - 1].
\end{aligned}$$

In addition to the safety constraints, a minimum value v_{ctrlmin} for the speed limits could be imposed:

$$v_{\text{ctrl},m,i}(l) \geq v_{\text{ctrlmin}} \quad \text{for all } (m, i) \in I_{\text{ctrl}} \text{ and } l \in [k, \dots, k + N_c - 1],$$

In practice, the variable speed limit signs display speed limits in increments of, e.g., 10 or 20 km/h. Therefore, the controller should produce discrete control signals. This is expressed by the constraint

$$v_{\text{ctrl},m,i}(l) \in \mathcal{V}_{m,i} \quad \text{for all } (m, i) \in I_{\text{ctrl}} \text{ and } l \in [k, \dots, k + N_c - 1], \quad (6)$$

where $\mathcal{V}_{m,i}$ is the set of discrete speed limit values in segment i of link m .

IV. A BENCHMARK PROBLEM

In order to illustrate the control framework presented above, we will apply it now to a benchmark problem.

A. Set-up

The benchmark set-up consists of a one origin, one freeway link of 12 km, and one destination, similar to Figure 3. The mainstream origin has two lanes with a capacity of 2000 veh/h each. The freeway link has two lanes, and consists of twelve segments of 1 km each. Segments 1 up to 5 and segment 12 are uncontrolled, while segments 6 up to 11 are equipped with a variable message sign where speed limits can be set. We use the same network parameters as in [25]: $T = 10$ s, $\tau = 18$ s, $\kappa = 40$ veh/lane/km, $\rho_{\text{max}} = 180$ veh/lane/km, $\rho_{\text{crit}} = 33.5$ veh/lane/km, $a_m = 1.867$ and $v_{\text{free}} = 102$ km/h.

Furthermore, we take $\eta_{\text{high}} = 65$ km²/h, $\eta_{\text{low}} = 30$ km²/h, $\alpha = 0.05$ and $\alpha_{\text{speed}} = 2$. Note that these

parameter values (and the speed limit and anticipation model) are not validated yet against real data.⁷

For the variable speed limits we have assumed that they can change once per minute, and that they cannot be less than $v_{\text{ctrlmin}} = 50$ km/h. This is imposed as a hard constraint in the optimization problem. For the cases where there is a safety constraint present, $v_{\text{maxdiff}} = 10$ km/h.

The input of the system is the traffic demand at the upstream end of the link and the (virtual) downstream density at the downstream end of the link. The traffic demand (inflow) has a constant value of 3900 veh/h, close to capacity (4000 veh/h). The downstream density equals the steady-state value of 28 veh/km, except for the pulse that represents the shock wave. The pulse was chosen large enough to cause a back-propagating wave in the segments, see Figures 5 and 6 (top). It is assumed that the upstream demand and downstream density is known, or predicted by an external algorithm. In practice, a combination of traffic measurements outside the controlled area and historical data could be used for prediction.

For the above scenario N_p and N_c will be tuned, and the performance (TTS) of the continuous and discrete controls with or without safety constraints are examined. In the discrete control case, the control values $v_{\text{ctrl},m,i}$ are in the set $\mathcal{V}_{m,i} = \mathcal{V} = \{50, 60, 70, 80, 90, 100, 110\}$ for all $(m, i) \in I_{\text{ctrl}}$.

The solution of the continuous speed control problem is calculated by the Matlab implementation of the SQP (sequential quadratic programming) algorithm `fmincon`. To get the discrete control signal the continuous signal is a rounded to a value in \mathcal{V} . Three different types of rounding are examined: The first (“round”) rounds the continuous control values to the nearest discrete value in \mathcal{V} , the second (“ceil”) rounds upwards to the nearest discrete value in \mathcal{V} , and the third (“floor”) rounds downwards to the nearest discrete value in \mathcal{V} .

This method of obtaining discrete control signals is heuristic but fast. It is also possible to use discrete optimization techniques such as tabu search, simulated annealing, or genetic algorithms, but since (as we will see) for this set-up and input the discretization method results in a performance that is comparable to that of the continuous version, it is not necessary to do so.

Note that it is not difficult to prove that the result of all of the three types of rounding will satisfy the safety constraints if the continuous signal satisfies them and if

⁷Validation will be performed in the future based on the data of the DYVERS experiment [8] where the speed limits were strictly enforced, and this data is expected to give a good indication of the achievable flow limitation of dynamic speed limits. Nevertheless, the most relevant effect, namely the capacity drop, is reproduced.

$v_{\max\text{diff}}$ is a multiple of the discretization step of the speed limits (here: multiple of 10 km/h). Since it does not make much sense to set $v_{\max\text{diff}}$ to another value than a multiple of the discretization step, this condition should not be a limitation.

The rolling horizon strategy is now implemented as follows. After the discretization, the first sample of the control signal is applied to the traffic system, and then the optimization–discretization steps are repeated. Note that this approach does not yield the same evolution and control signals as an approach in which first the continuous signal is computed (using the rolling horizon approach) for the entire simulation period at once, rounded, and then applied for the whole simulation period. This is because in the first approach the different traffic behavior caused by the discretization is already taken into account in the each subsequent MPC iteration.

In the next section we will compare the performance of the discrete control to the performance achieved by the continuous valued control without constraints, and the effect of introducing the safety constraints is examined.

B. Results

The results of the simulations of the no-control and the control case with continuous speed limits without constraints is displayed in Figure 6 (top and middle). Since high densities correspond to low speeds and low densities to high speeds, the corresponding speeds plots (not shown here) have roughly the same shape as the plots in Figure 6, but with the vertical axis flipped.

The propagation speed of the shock wave in the uncontrolled simulation is around 22 km/h (11 km/30 min), which is somewhat higher than the 15–20 km/h range reported in the literature⁸ [34], [21].

In the controlled case the shock wave disappears after approximately 90 minutes, while in the no-control case, the shock wave travels through the whole link. The speed limits are active in segments 6 up to 10; the speed limit in segment 11 has higher values than the critical speed and is not limiting the flow (see Figure 8). The active speed limits start to limit the flow at $t = 5$ min and create a low-density wave traveling downstream, see the small dip in Figure 6 (middle and bottom). This low-density wave meets the shock wave and reduces its density just enough to stop it. The tail of the shock wave has a fixed

⁸This is probably the consequence of the chosen model parameters that are not calibrated in this study (e.g., η_{high} and η_{low}). The necessity of calibration is a topic for future research. What is relevant in this paper is the occurrence of the capacity-drop: the outflow of the shock wave is less than the freeway capacity. Therefore, resolving the shock wave will result in a higher flow and an improvement of the total time spent.

location while the head dissolves into free flow traffic, and the shock wave eventually dissolves completely.

The speed limits persist until the shock wave (to be precise, the high-density region) is completely dissolved. The speed limits in Figure 8 start to increase after $t = 35$ min and return gradually to a high value that is not limiting the flow anymore. Since the shock wave is completely dissolved within the freeway stretch, traffic upstream the stretch will not be influenced by the control. Note, however, that in the uncontrolled situation the shock wave does propagate upstream and will impede the upstream traffic.

Note that the speed limit values after 90 to 120 minutes are in the range of 70–100 km/h. These values do not limit the flow because they are higher than the critical speed (60 km/h in our case), and neither do they limit the traffic speed, because the traffic speeds (not shown here) were lower than the displayed speed limits.

The TTS was 1835.3 veh.h in the no-control case and 1466.7 veh.h in the controlled (continuous, unconstrained) case, which is an improvement of 20.1%.

The relative improvement of the performance as a function of N_p and N_c is shown in Figure 7. The performance depends stronger on N_p , but for $N_p \geq 10$ (which corresponds to 10 min, and is somewhat larger than the maximum travel time from segment 6 to the exit of the link) the graphs become nearly flat. We select $N_p = 10$ and $N_c = 8$ for the further analysis.

The result of the several types of discretization is shown in Table I for the simulations without safety constraints and in Table II for the simulations with safety constraints. The performance loss caused by the discretization is small in the “round” and “ceil” cases, but large for “floor”. The cause for this performance degradation in the latter case is probably that rounding the speed limits downwards limits the flow too much. This is also in accordance with the fact that the performance of “floor” is better in the constrained case than in the unconstrained case, because the constraints prevent too low speed limits, so the downwards rounded version will on the average also be higher, resulting in a higher flow. In the other cases the inclusion of the constraints result in a small performance loss, which is in accordance with the general expectation that the introduction of extra constraints usually results in lower performance.

The performance improvement for $N_p = 10$, $N_c = 8$ in the constrained “ceil” case is 17.3%, which is close to the improvement of the unconstrained “ceil” case (18.3%), and comparable to the improvement of 20.1% in the unconstrained continuous case. Figure 8 (bottom) shows the values of the optimal speed limits in the discrete (“ceil”) case with safety constraints and

Table I

THE RELATIVE IMPROVEMENT OF THE PERFORMANCE (TTS) FOR SEVERAL COMBINATIONS OF N_p AND N_c , AND FOR THE CONTINUOUS SPEED LIMITS AND THE THREE DISCRETE SPEED LIMITS: "ROUND", "CEIL", AND "FLOOR"; WITHOUT SAFETY CONSTRAINTS.

Horizon		Relative improvement (%)			
N_p	N_c	continuous	round	ceil	floor
9	4	19.6	17.5	17.9	-2.2
9	6	19.6	19.1	18.9	3.9
9	8	19.8	15.0	17.6	6.9
10	4	19.9	17.9	19.6	-1.1
10	6	20.0	19.6	19.3	2.9
10	8	20.1	15.2	18.3	5.9
11	4	20.0	18.0	19.8	-1.1
11	6	20.0	17.7	19.8	1.3
11	8	20.0	19.9	19.4	5.5
12	4	20.1	15.5	20.0	-2.2
12	6	20.1	19.7	20.0	1.3
12	8	20.2	19.8	20.0	5.7

Table II

THE RELATIVE IMPROVEMENT OF THE PERFORMANCE (TTS) FOR SEVERAL COMBINATIONS OF N_p AND N_c , AND FOR THE CONTINUOUS SPEED LIMITS AND THE THREE DISCRETE SPEED LIMITS: "ROUND", "CEIL", AND "FLOOR"; WITH SAFETY CONSTRAINTS.

Horizon		Relative improvement (%)			
N_p	N_c	continuous	round	ceil	floor
9	4	19.4	16.4	18.0	0.2
9	6	19.5	19.3	19.0	12.3
9	8	19.4	18.4	11.4	11.9
10	4	19.5	15.5	18.5	1.4
10	6	19.6	19.4	18.0	9.0
10	8	19.7	19.1	17.2	11.0
11	4	19.6	15.4	18.2	0.4
11	6	19.7	19.8	19.6	7.3
11	8	19.9	19.7	19.3	5.5
12	4	19.7	14.7	19.3	1.8
12	6	19.9	19.9	19.7	12.5
12	8	19.9	19.3	19.6	13.4

$$N_p = 10, N_c = 8.$$

The computation time varied between 3 and 25 min on a 500 MHz Pentium III PC, which is at least four times faster than real time. It is expected that the computation time varies linearly with the number of segments and the length of the prediction horizon, and exponentially with the number of control inputs and the control horizon.

The fact that the same traffic flow model was used as a prediction model and as the model for the controlled traffic system may have influenced (probably positively) the achievable improvement of the TTS. Nevertheless, the presented simulations give an indication that shock waves can be resolved by dynamic flow-limiting speed limits, and that a significant improvement of traffic flow and travel times can be achieved. In case there is a mismatch between the process model and the controller

model several techniques exist to ensure robustness [35].

V. CONCLUSIONS AND FUTURE RESEARCH

We have applied model predictive control (MPC) to optimally coordinate variable speed limits with the aim of suppressing shock waves. The purpose of the control was to find the control signals that minimize the total time that vehicles spend in the network.

We have applied the developed control framework to a benchmark network consisting of a link of 12 km, where 6 segments of 1 km are controlled by speed limits. It was shown that coordinated control with continuous speed limits (base case) is effective against shock waves.

The performance loss caused by discrete speed limits and the inclusion of safety constraints was examined. The performance of the discrete safety-constrained speed limits was comparable that of the base case if the discrete speed limits were generated by "round" or "ceil". In all of these cases the coordination of speed limits eliminated the shock wave entering from the downstream end of the link. The coordinated case resulted in a network where the outflow was sooner restored to capacity, and in an improvement (decrease) of the total time spent of 17.3 %.

Topics for further research include:

- Comparison of the discrete MPC approach with other existing approaches;
- Further examination of the trade-off between efficiency and optimality for rounding versus full discrete optimization;
- Study of a real freeway stretch, including model calibration with real data; the examination of the necessity of on-line calibration; simulation of other set-ups and scenarios; validation of the new modeling assumptions regarding the speed limits and the main stream origin;
- Further investigation of the effectiveness of MPC for optimal coordination of speed limits for a wider range of scenarios, networks, traffic flow models and/or model parameters; explicit inclusion of modeling errors, different internal and external models and unpredicted disturbances (demands);
- Further study of the capacity drop and metastability phenomena;
- Inclusion of extra control measures in addition to speed limits (such as ramp metering, dynamic lane assignment, route guidance, reversible lanes, etc.).
- Investigation of issues relevant for real-life application, such as the extension of the control system with a traffic state estimation (for which a Kalman filtering approach could be followed similarly to [36]) and a traffic demand prediction module, investigation of the robustness of the control

approach against unknown disturbances, state estimation errors, and a model mismatch; investigation of the requirements of the physical lay-out of the freeway stretch, and the estimation of the expected improvement for a real scenario.

REFERENCES

- [1] M. Papageorgiou, J. Blosseville, and H. Hadj-Salem, "La fluidification des roades de l'Île de France: Un projet d'importance." Dynamic Systems and Simulation Laboratory, Technical University of Crete, Chania, Greece, Tech. Rep., 1998, internal Report No. 1998-17.
- [2] M. J. Cassidy and R. L. Bertini, "Some traffic features at freeway bottlenecks," *Transportation Research Part B*, vol. B33, pp. 25–42, 1999.
- [3] A. Hegyi, "Model predictive control for integrating traffic control measures," Ph.D. thesis, TRAIL Thesis Series T2004/2, Delft University of Technology, Delft, The Netherlands, Feb. 2004, ISBN 90-5584-053-X, http://www.dsc.tudelft.nl/~ahegy/thesis/dep_hegyi_20040203.pdf.
- [4] H. Zackor, "Self-sufficient control of speed on freeways," in *Proceedings of the International Symposium on Traffic Control Systems*, vol. 2A. Berkeley, California: California University, Aug. 6–9 1979, pp. 226–249.
- [5] S. Smulders, "Control by variable speed signs — the Dutch experiment," in *Proceedings of the Sixth International Conference on Road Traffic Monitoring and Control*, ser. IEE Conference Publication. London: IEE, Apr. 28–30 1992, pp. 99–103.
- [6] E. v. d. Hoogen and S. Smulders, "Control by variable speed signs: results of the dutch experiment," in *7th International Conference on Road Traffic Monitoring and Control*, ser. IEE Conference Publication No. 391, London, England, Apr. 26–28, 1994, pp. 145–149.
- [7] E. J. Hardman, "Motorway speed control strategies using SISTM," in *Road Traffic Monitoring and Control*, ser. Conference Publication, no. 422. IEE, Apr. 23–25 1996, pp. 169–172.
- [8] DHV Milieu en Infrastructuur BV, "Eindevaluatie SlimRijden-verkeerskundige- en weggebruikersevaluatie (Final evaluation of the project SlimRijden (SmartDriving) – from a traffic engineering and a road users' point of view)," AVV Traffic Research Centre, Dutch Ministry of Transport, Public Works and Water Management, Tech. Rep., 2003, in Dutch.
- [9] C.-C. Chien, Y. Zhang, and P. A. Ioannou, "Traffic density control for automated highway systems," *Automatica*, vol. 33, no. 7, pp. 1273–1285, 1997.
- [10] H. Lenz, R. Sollacher, and M. Lang, "Nonlinear speed-control for a continuum theory of traffic flow," in *14th World Congress of IFAC*, vol. Q, Beijing, China, 1999, pp. 67–72.
- [11] —, "Standing waves and the influence of speed limits," in *Proceedings of the European Control Conference 2001*, Porto, Portugal, 2001, pp. 1228–1232.
- [12] S. Smulders, "Control of freeway traffic flow by variable speed signs," *Transportation Research Part B*, vol. 24B, no. 2, pp. 111–132, 1990.
- [13] Rijkswaterstaat Directie Zuid-Holland, "Evaluatie 80 km/uurmaatregel A13 Overschie, Doorstroming and Verkeersveiligheid (Evaluation of the 80 km/h measure on the A13 at Overschie, Traffic Flow and Safety)," Ministry of Transport, Public Works and Water Management, Rijkswaterstaat, Directie Zuid-Holland - VIV, Tech. Rep., 2003, in Dutch.
- [14] J. Wesseling, K. Hollander, S. Teeuwisse, M. Keuken, H. Spoelstra, R. Gense, E. Burgwal, L. Hermans, J. Voerman, P. Kumm, and J. Elshout, "Onderzoek naar de effecten van de 80 km/h-maatregel voor de A13 op de luchtkwaliteit in Overschie (Investigation of the effects of the 80 km/h measure at the A13 on the air quality in Overschie)," TNO (on behalf of Rijkswaterstaat, Directie Zuid-Holland), Tech. Rep., 2003, in Dutch.
- [15] J. K. Wilkie, "Using variable speed limit signs to mitigate speed differentials upstream of reduced flow locations," Department of Civil Engineering, Texas A& M University, College Station, Texas 77843, Tech. Rep., Aug. 1997, prepared for CVEN 677 Advanced Surface Transportation Systems.
- [16] P. Y. Li, R. Horowitz, L. Alvarez, J. Frankel, and A. M. Robertson, "Traffic flow stabilization," in *Proceedings of the American Control Conference*, Seattle, Washington, June 1995, pp. 144–149.
- [17] A. Alessandri, A. Di Febbraro, A. Ferrara, and E. Punta, "Nonlinear optimization for freeway control using variable-speed signaling," *IEEE Transactions on Vehicular Technology*, vol. 48, no. 6, pp. 2042–2052, Nov. 1999.
- [18] A. Di Febbraro, T. Parisini, S. Sacone, and R. Zoppoli, "Neural approximations for feedback optimal control of freeway systems," *IEEE Transactions on Vehicular Technology*, vol. 50, no. 1, pp. 302–312, Jan. 2001.
- [19] P. Breton, A. Hegyi, B. De Schutter, and H. Hellendoorn, "Shock wave elimination/reduction by optimal coordination of variable speed limits," in *Proceedings of the IEEE 5th International Conference on Intelligent Transportation Systems*, Singapore, Sept. 3–6 2002, pp. 225–230.
- [20] A. Hegyi, B. De Schutter, and H. Hellendoorn, "Optimal coordination of variable speed limits to suppress shock waves," in *Proceedings of the 82nd Annual Meeting of the Transportation Research Board*. Transportation Research Board, Jan. 2003.
- [21] B. S. Kerner and H. Rehborn, "Experimental features and characteristics of traffic jams," *Physical Review E*, vol. 53, no. 2, pp. R1297–R1300, February 1996.
- [22] E. Camacho and C. Bordons, *Model Predictive Control in the Process Industry*. Berlin, Germany: Springer-Verlag, 1995.
- [23] J. Maciejowski, *Predictive Control with Constraints*. Harlow, England: Prentice Hall, 2002.
- [24] F. Allgöwer, T. Badgwell, J. Qin, J. Rawlings, and S. Wright, "Nonlinear predictive control and moving horizon estimation – An introductory overview," in *Advances in Control: Highlights of ECC '99*, P. Frank, Ed. Springer, 1999, pp. 391–449.
- [25] A. Kotsialos, M. Papageorgiou, and A. Messmer, "Optimal coordinated and integrated motorway network traffic control," in *Proceedings of the 14th International Symposium on Transportation and Traffic Theory (ISTTT)*, Jerusalem, Israel, July 1999, pp. 621–644.
- [26] M. Papageorgiou, J.-M. Blosseville, and H. Haj-Salem, "Modelling and real-time control of traffic flow on the southern part of Boulevard Périphérique in Paris: Part II: coordinated on-ramp metering," *Transportation Research Part A*, vol. 24A, no. 5, pp. 361–370, 1990.
- [27] Technical University of Crete and A. Messmer, *METANET – A simulation program for motorway networks*, Technical University of Crete, Dynamic Systems and Simulation Laboratory and A. Messmer, Nov. 2001.
- [28] M. Papageorgiou, J.-M. Blosseville, and H. Hadj-Salem, "Modelling and real-time control of traffic flow on the southern part of Boulevard Périphérique in Paris: Part I: Modelling," *Transportation Research Part A*, vol. 24A, no. 5, pp. 345–359, 1990.
- [29] Rand EUROPE, "Validatie METANET, een rapport voor Rijkswaterstaat AVV (Validation METANET, a report for Directorate-General for Public Works and Water Management, Transport Research Centre)," Rand EUROPE, Leiden, The Netherlands, Tech. Rep. 0061, July 2001, in Dutch.

- [30] D. Ngoduy and S. P. Hoogendoorn, "An automated calibration procedure for macroscopic traffic flow models," in *Proceedings of the 10th IFAC Symposium on Control in Transportation Systems (CTS 2003)*, Tokyo, Japan, Aug. 2003, pp. 295–300.
- [31] M. Cremer, *Der Verkehrsfluss auf Schnellstrassen (Traffic flow on freeways)*, ser. Fachberichte Messen, Steuern, Regeln. Berlin: Springer-Verlag, 1979, vol. 3, in German.
- [32] S. Hoogendoorn, "Multiclass continuum modelling of multiclass traffic flow," Ph.D. dissertation, Delft University of Technology, TRAIL Thesis Series, Delft, The Netherlands, Dec. 1999.
- [33] A. Hegyi, B. De Schutter, and J. Hellendoorn, "Optimal coordination of variable speed limits to suppress shock waves," *Transportation Research Record*, no. 1852, pp. 167–174, 2004.
- [34] B. S. Kerner, "Empirical features of congested patterns at highway bottlenecks," in *Proceedings of the 81st Annual Meeting of the Transportation Research Board*, Washington, D.C., 2002.
- [35] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, pp. 789–814, 2000.
- [36] Y. Wang, M. Papageorgiou, and A. Messmer, "RENAISSANCE: A real-time motorway network traffic surveillance tool," in *Preprints 10th IFAC Symposium on Control in Transportation Systems*, Tokyo, Japan, Aug. 4–6 2003, pp. 235–240.

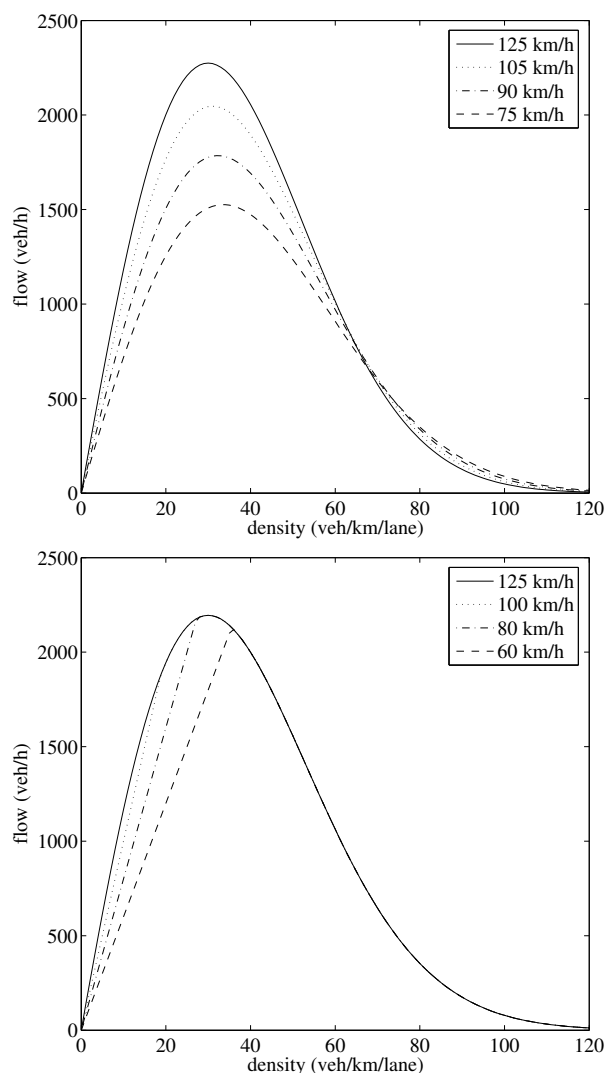


Figure 4. Illustrations of the discussed speed limit models. Slopes correspond to speeds in these figures. **Top:** In the model proposed by Cremer the flow-density relation is approximately scaled down linearly with the speed limit. This also influences traffic with a speed that is lower than the speed limit, which is unrealistic. **Bottom:** In the proposed speed limit model the speed limit changes only the shape of the fundamental diagram for densities where the speed would have been higher (or: the slope would have been steeper) than the applied speed limit. The straight lines represent the region where the speed limits do limit the traffic speed, for higher densities (corresponding to lower speeds than the speed limit) the new diagram coincides with the original one (solid line).

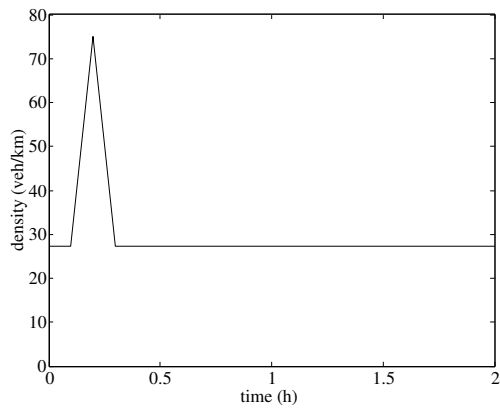


Figure 5. The downstream density scenario considered in the experiments.

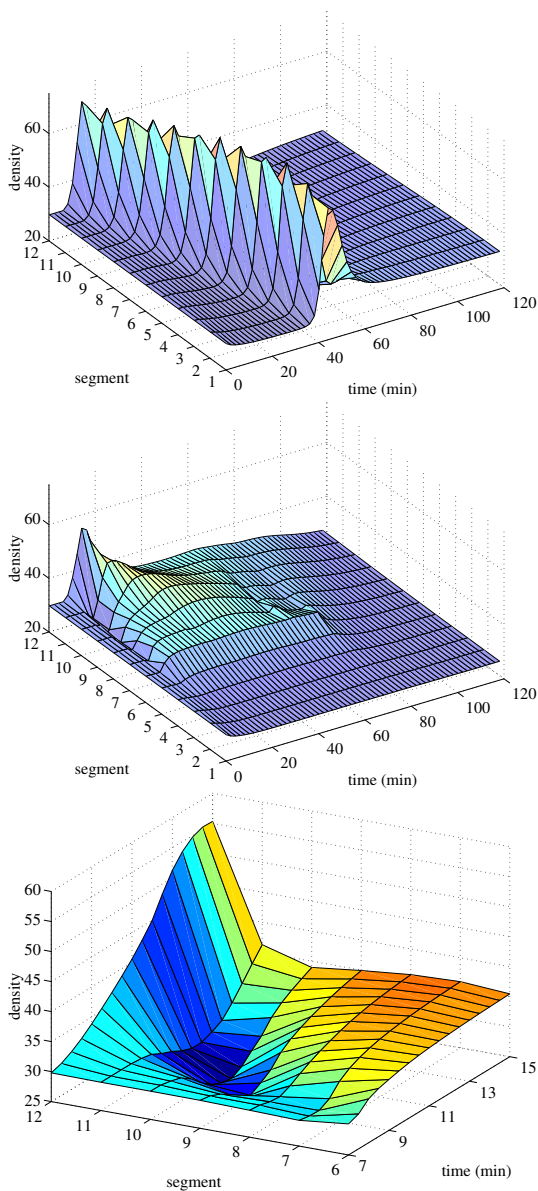


Figure 6. The shock wave propagates through the link in the no-control case (top). In the coordinated control case, the shock wave disappears after approximately 90 minutes (middle), and the zoom-in on the dip for the coordinated control case (bottom).

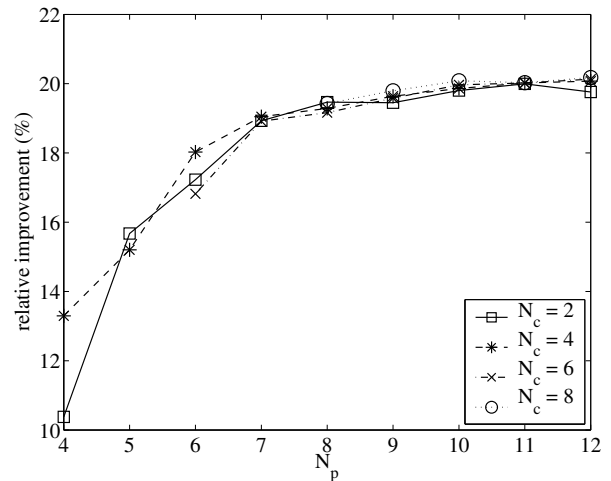


Figure 7. The relative improvement of the performance (Total Time Spent) in the continuous, unconstrained case compared to the no-control case as a function of N_p for several values of N_c . The sensitivity to N_p is much higher than that to N_c .

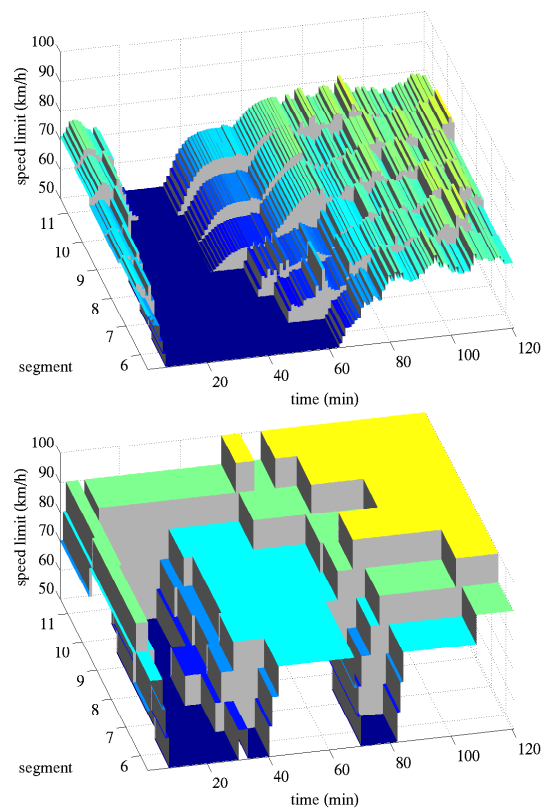


Figure 8. The speed limits for the continuous case without safety constraints and $N_p = 10$, $N_c = 8$ (top). The speed limits for the discrete ("ceil") case with safety constraints and $N_p = 10$, $N_c = 8$ (bottom).