

Technical report CSE02-014

MPC-based optimal coordination of variable speed limits to suppress shock waves in freeway traffic*

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, "MPC-based optimal coordination of variable speed limits to suppress shock waves in freeway traffic," *Proceedings of the 2003 American Control Conference*, Denver, Colorado, pp. 4083–4088, June 2003.

Control Systems Engineering
Faculty of Information Technology and Systems
Delft University of Technology
Delft, The Netherlands
phone: +31-15-278.51.19 (secretary)
fax: +31-15-278.66.79
Current URL: <http://www.dcsc.tudelft.nl>

*This report can also be downloaded via http://pub.deschutter.info/abs/02_014.html

MPC-based optimal coordination of variable speed limits to suppress shock waves in freeway traffic

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

Control Systems Engineering, Fac. ITS, Delft University of Technology
P.O. Box 5031, 2600 GA Delft, The Netherlands, {a.hegyi,b.deschutter,j.hellendoorn}@its.tudelft.nl

Abstract

At freeway bottlenecks 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 and dangerous situations. Dynamic speed limits can be used to eliminate or at least to reduce the effects of shock waves. However, in order to prevent the occurrence of new shock waves and/or negative impacts on the traffic flows in other locations, coordination of the variable speed limits is necessary. In this paper we further extend our results in connection with a model predictive control (MPC) approach to optimally coordinate variable speed limits for freeway traffic. First of all, 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-valued speed limits.

1 Introduction

When using dynamic speed limits for congestion reduction two aspects are relevant: homogenization and prevention of breakdown. The basic idea behind homogenization [1, 4, 15, 13] is that speed limits reduce the speed differences between vehicles, by which a higher (and safer) flow can be achieved. The homogenization approach typically uses speed limits that are close to the critical speed (i.e. the speed that corresponds to the capacity flow; see Figure 1). The results in [15] indicate that the effect of homogenization on freeway performance is small; however, a positive safety effect can be expected. The traffic breakdown prevention approach [3, 7] focuses more on preventing too high densities. We consider the use of variable speed limits to prevent traffic breakdown.

In the literature several methods are described to synthesize suitable control laws for speed limit control, such as multi-layer control [10], sliding-mode control [7], or optimal control [1, 4]. This paper extends the results of our previous paper [2], in which we have already demonstrated the effectiveness of continuous-valued 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-valued speed limits.

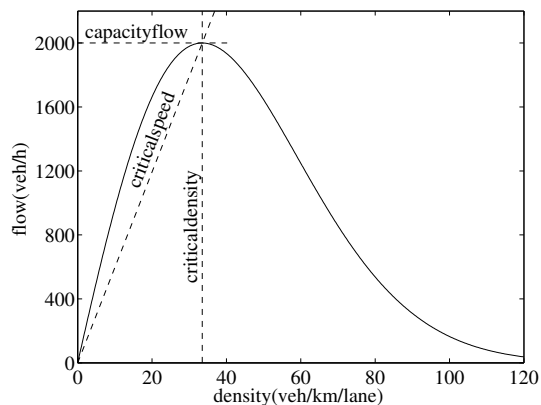


Figure 1: A typical example of the fundamental diagram, which represents the relation between density and flow for a given freeway section.

2 Problem statement

2.1 Shock waves in traffic flows

It is well known [5] that some types of traffic jams move upstream with approximately 15 km/h. These moving jams are called waves. As they can remain existent for a long time, every vehicle that enters the freeway upstream of the congested 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 waves¹ speed limits can be used as follows. On some sections upstream of a shock wave speed limits are imposed and consequently the inflow of the congested area is reduced. When the inflow of the jammed area is smaller than its outflow, the jam will eventually dissolve. In other words, the speed limits create a low density wave (with a density lower than it would be in the uncontrolled situation) that propagates downstream. This low density wave meets the shock wave and compensates its high density, which reduces or eliminates the shock wave. A point of criticism could be that the above approach reduces the shock wave, but at the cost of creating new shock waves upstream of the sections controlled by speed limits. However, if the speed limits are optimized properly,

¹Lighthill and Whitham [8] introduced the term shock wave for waves formed by several waves running together. At the shock wave fairly large reductions in velocity occur very quickly. We use the term “shock wave” for any wave (the moving congested areas) and we do not distinguish between waves and shock waves, because in practice any wave is undesired.

they will never create a shock wave that gives rise to higher delays than in the uncontrolled case. This issue will be addressed by the coordinated MPC-based approach for variable speed limit control presented in this paper.

2.2 Coordination and prediction

In practice, dynamic traffic management often still operates based on local data only. However, considering the effect of the measures on the network level has in general many advantages compared to local control. Hence, a network-wide *coordination* of the control measures based on global data is certainly useful. The coordination of the control signals is obtained by reformulating the control design problem over a given time horizon as an optimization problem that yields the optimal speed limit settings (see Section 3).

Since we want to determine control signal settings that are optimal for a given freeway network and since the effect of a given control measure on more distant parts of the network might only be visible/measurable after some time, an accurate *prediction* of the future evolution of the traffic flows in the network is also necessary. In particular, 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 can not be observed instantaneously,² so the prediction should be made at least up to the point where the improvement can be observed. For the prediction we will use a modified and extended version of the METANET traffic flow model.

Besides prediction and coordination the speed limit control problem has some other characteristics that impose specific requirements to the control strategy:

- 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 [11] 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%. This effect can be explained by the fact that the number of vehicles in the network is equal to the accumulated net inflow of the network. But the outflow is lower when the traffic is congested³, so the queue grows faster, and consequently congestion will last longer, and the outflow will be low for a longer time (the time that the queue needs to dissolve). So 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. Furthermore, because this flow improvement is relatively small, and there are always disturbances present in the traffic flow *feedback control* is required. In this way imprecisions of control and traffic disturbances can be observed and appropriate control actions can be taken.

²We will see that the speed limits have to slow down a part of the traffic first in order to dissolve the shock wave.

³The congestion after a breakdown usually has an outflow that is (only 5–10%!) lower than the available capacity; this is the so called capacity-drop phenomenon.

- As the speed limit signs used in practice display speed limits in increments of, e.g., 10 or 20 km/h, the controller should produce discrete-valued control signals.
- For safety it is often required that the driver should not encounter a decrease in the speed limit larger than a pre-specified amount. The controller should be able to take this kind of constraints into account.

The control strategy presented in the next section addresses these issues.

3 MPC-based coordination of speed limits

3.1 Model predictive control

We will use a model predictive control (MPC) scheme to solve the problem of optimal coordination of speed limits. In MPC, at each time step k the optimal control signal is computed (by numerical optimization) over a prediction horizon N_p . In order to reduce the number of variables and to improve the stability of the system, in MPC a control horizon N_c ($< N_p$) is selected. After the control horizon has been passed the control signal is usually taken to be constant. In addition, a rolling horizon strategy is used, which means that at each time step only the first sample of the optimal control signal is applied to the system; afterwards the time axis is shifted one sample step, the model is updated, and the procedure is restarted. The rolling horizon approach results in an on-line predictive and adaptive control scheme that allows us to take changes in the system or in the system parameters into account by regularly updating the model of the system or the predicted demands as new measurements from the traffic sensors become available. For more information on MPC we refer the interested reader to [9, 14] and the references therein.

3.2 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 a slightly modified version of the METANET model [6, 12]. The modifications are introduced for better modeling of shock waves and the effect of speed limits. Note that the MPC approach is generic and will find the optimal speed limits independent from the model that is used (e.g., the way that speed limits enter the model), so the modifications do not interfere in any way with the effectivity of MPC. For the sake of brevity, we will describe only those parts of the METANET model that are relevant for interpreting and understanding the simulation results of our benchmark network (see Section 4).

3.2.1 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. Where major changes occur in the characteristics of the link or in the road geometry (on/off-ramp), a node is placed. Each link m is divided into N_m segments of length L_m (see Figure 2). Each

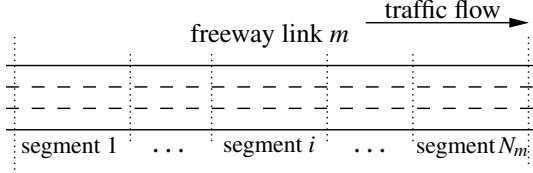


Figure 2: In the METANET model a freeway link is divided into segments.

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 outflow of each segment is given by

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

where λ_m denotes the number of lanes of segment m . The principle of conservation of vehicles yields:

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

The mean speed depends on the previous speed plus a relaxation term, a convection term, and an anticipation term:

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

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 $v_{\text{free},m}$ the free-flow speed, and $\rho_{\text{crit},m}$ the critical density.

Origins are modeled with a simple queue model:

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

with $w_o(k)$ the queue length for origin o , $d_o(k)$ the demand, and the outflow $q_o(k)$ given by

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

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

3.2.2 Extensions to the METANET model: Since the original METANET model does not describe the effect of speed limits, we have slightly modified the equation for the desired speed (3) to incorporate speed limits. The second extension regards the modeling of the different nature of

a mainstream origin as opposed to an on-ramp origin. The third extension considers the different effect of the downstream density gradient on the speed (cf. the anticipation term in (2)) when this gradient is positive or negative.

In some publications the effect of the speed limit is expressed by scaling down the desired speed-density diagram [1]. 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 even when the vehicles are traveling at 80 km/h. Furthermore, scaling down the desired speed also reduces 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 1). These 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 desired speed based on the experienced density, and the desired speed caused by the variable speed limit v_{ctrl} :

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

To express the different nature of a mainstream origin link o compared to 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), we use a modified version of (4) with another flow constraint, 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 on 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 on the first segment whichever is smaller. So if o is the origin of 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 :

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

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

Since the effect of a higher downstream density is usually stronger than the effect of a lower downstream density, we distinguish between these two cases. The sensitivity of the speed to the downstream density is expressed by parameter η . In (2) η is a global parameter that 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:

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

3.3 Constraints

For the safe operation of a speed control system, it is required that the maximum decrease of speed limits that a driver can encounter (v_{maxdiff}) is limited. There are 3 situations where a driver can encounter a different speed limit value: (1) when the speed limit changes in the next time step on 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 at that time step. The maximum speed difference constraints in these 3 situations are formulated as follows:

$$v_{\text{ctrl},m,i}(l-1) - v_{\text{ctrl},m,i}(l) \leq v_{\text{maxdiff}} \quad \text{for all } (m,i) \in I_{\text{speed}}$$

$$v_{\text{ctrl},m,i}(l) - v_{\text{ctrl},m,i+1}(l) \leq v_{\text{maxdiff}} \\ \text{for all } (m,i) \in I_{\text{speed}} \text{ such that } (m,i+1) \in I_{\text{speed}},$$

$$v_{\text{ctrl},m,i}(l-1) - v_{\text{ctrl},m,i+1}(l) \leq v_{\text{maxdiff}} \\ \text{for all } (m,i) \in I_{\text{speed}} \text{ such that } (m,i+1) \in I_{\text{speed}},$$

for $l \in [k, \dots, k + N_c]$, and with I_{speed} the set of pairs of indices (m,i) of the links and segments where speed control is applied. In addition to the above safety constraints, the speed limits are often subject to a minimum value $v_{\text{ctrl},\text{min}}$:

$$v_{\text{ctrl},m,i}(l) \geq v_{\text{ctrl},\text{min}} \quad \text{for all } (m,i) \in I_{\text{speed}} \text{ and } l \in L(k).$$

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-valued control signals. This is expressed by the constraint

$$v_{\text{ctrl},m,i}(l) \in \mathcal{V}_{m,i} \quad \text{for all } (m,i) \in I_{\text{speed}} \text{ and } l \in L(k), \quad (5)$$

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

3.4 Objective function

We consider the following objective function:

$$J(k) = T \sum_{l=k}^{k+N_p-1} \left\{ \sum_{(m,i) \in I_{\text{all}}} \rho_{m,i}(l) L_m \lambda_m + \sum_{o \in O_{\text{all}}} w_o(l) \right\} + \\ a_{\text{speed}} \sum_{l=k}^{k+N_c-1} \sum_{(m,i) \in I_{\text{speed}}} \left(\frac{v_{\text{ctrl},m,i}(l) - v_{\text{ctrl},m,i}(l-1)}{v_{\text{free},m}} \right)^2,$$

where I_{all} and O_{all} are the sets of indices of all pairs of segments and links and of all origins respectively. 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 a_{speed} .

3.5 MPC-based speed limit control

3.5.1 Computing the optimal controls: Now the optimal control signal can be computed by minimizing the objective function $J(k)$ subject to the model of the system (which can be considered as a system of equality constraints) and subject to the safety constraints.

Let us first consider the case with real-valued control signals, i.e., without condition (5). This results in a nonlinear, nonconvex optimization problem, which can be solved using, e.g., an SQP approach. Note that due to the rolling horizon approach we can use the current optimal control signal (i.e., for time step k) as a good initial guess for the next optimization (i.e., for time step $k+1$).

If discrete-valued control signals are considered, nonlinear integer programming is required. However, due to the inherent complexity of integer optimization problems and due to the large number of variables, such an approach is not tractable in practice, especially as repeated on-line optimization is required in MPC. Hence, we propose to use a rounded version of the continuous optimization result for the discrete-valued control signals. In particular, for the benchmark problem of Section 4 we will examine three different types of discretization: the first (“round”) rounds the continuous control values to the nearest discrete value, the second (“ceil”) to the nearest discrete value that is higher than the continuous value, and the third (“floor”) to the nearest discrete value that is lower than the continuous value.

After the discretization the first sample of the control signal is applied to the traffic system and then the optimization–discretization steps are repeated. It is important to note that this way of rounding is not the same as rounding the continuous signal of the whole prediction horizon at once, because here the different traffic behavior caused by the discretization is already taken into account in the next MPC iteration. The above 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 the discretization method is very fast and usually results in a performance that is comparable to the continuous version (cf. Section 4), the discretization method is often sufficient.

4 A benchmark problem

In order to illustrate the control framework presented above we will now apply it to benchmark set-up consisting of a freeway link equipped with variable speed signs.

4.1 Set-up

The benchmark set-up consists of an origin, a freeway link, and a destination, as in Figure 2. The mainstream origin O_1 has 2 lanes with a capacity of 2000 veh/h each. The freeway link L_1 has 2 lanes, is 12 km long, and consists of $N_1 = 12$ segments of 1 km each. Segments 1 up to 5

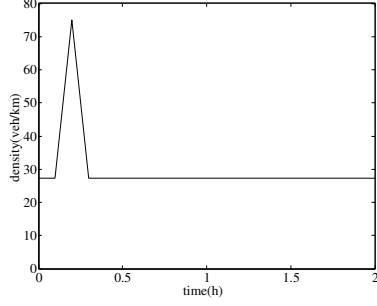


Figure 3: The downstream density scenario.

and 12 are uncontrolled, segments 6 up to 11 are equipped with a variable message sign where speed limits can be set. The choice for the 5 uncontrolled upstream segments was made to be sure that the boundary conditions do not play any dominant role. Link L_1 ends in destination D_1 . We use the same network parameters as in [6]: $T = 10$ s, $\tau = 18$ s, $\kappa = 40$ veh/lane/km, $\rho_{\max} = 180$ 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, and $a_{\text{speed}} = 2$. For the variable speed limits we have assumed that they can change only every minute, and that they cannot be less than $v_{\text{ctrl,min}} = 50$ km/h. This is imposed as a hard constraint in the optimization problem. If there is a safety constraint, then we take $v_{\text{maxdiff}} = 10$ 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 3850 veh/h, close to capacity (4000 veh/h). The downstream density equals the steady-state value of 27.4 veh/km, except for the pulse that represents the shock wave (see Figure 3). The pulse was chosen large enough to cause a backpropagating wave in the segments (see Figure 4 (top)). In the discrete-valued control case, the control values $v_{\text{ctrl},m,i}$ are in the set $\{50, 60, 70, 80, 90, 100, 110\}$.

4.2 Results

The results of the simulations of the no control and the control with continuous speed limits without constraints are displayed in Figure 4. In the controlled case the shock wave disappears after approximately 90 min, while in the no control case the shock wave travels through the whole link. The active speed limits start to limit the flow at $t = 5$ min and create a low density wave traveling downstream (the small dip in Figure 4 (bottom) and Figure 5). This low density wave meets the shock wave traveling upstream and reduces its density just enough to stop it. So, the tail of the shock wave has a fixed location while the head dissolves into free flow traffic as in the uncontrolled situation, which means that 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 6 (top) start to increase after $t = 35$ min and return gradually to a high value that is not limiting the flow anymore. The TTS was 1835.3 veh.hours in the no control case and 1466.7 veh.hours in the controlled (continuous, un-

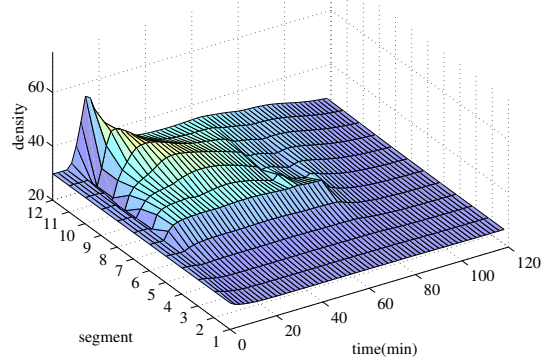
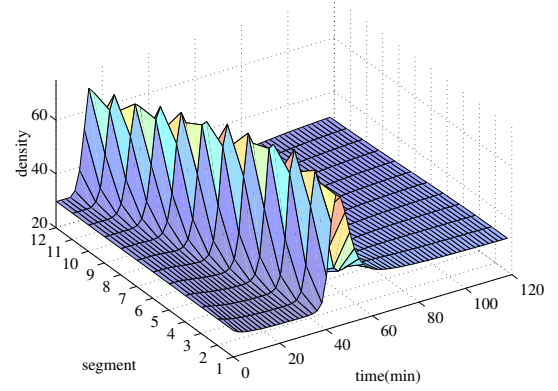


Figure 4: The shock wave propagates backwards through the link in the no control case (top). In the coordinated control case (bottom) the shock wave does not propagate through the entire link and completely disappears after approximately 90 min.

Relative improvement (%)			
continuous	round	ceil	floor
20.1	15.2	18.3	5.9

Table 1: The relative improvement of TTS for the continuous-valued speed limits and the three discrete-valued speed limits (without safety constraints).

constrained) case, which is an improvement of 20.1 %. We have chosen $N_p = 10$ and $N_c = 8$ (cf. tuning rules of [2]).

The results for the several types of discretization are shown in Table 1. The performance loss caused by the discretized speed limits is small in the “round” and “ceil” cases, but large for “floor”. If the safety constraints are included, the results are comparable to Table 1. The performance improvement in the constrained “ceil” case is 17.3 %, compared to 20.1 % in the unconstrained continuous case. Figure 6 (bottom) shows the values of the optimal speed limits discrete (“ceil”) case with safety constraints.

5 Conclusions and future research

We have presented a model predictive control approach to optimally coordinate variable speed limits. The purpose of the control was to minimize the total time that vehicles spend in the network. We have applied the developed control framework to a benchmark network. It was shown that coordinated control with continuous-valued speed lim-

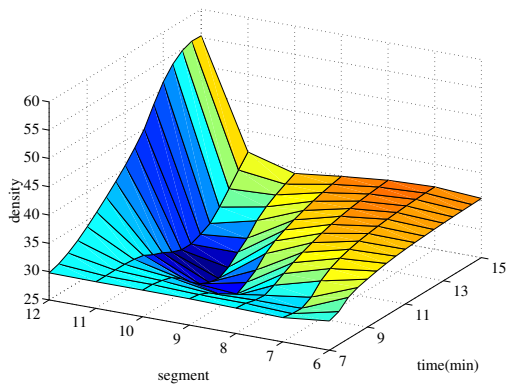


Figure 5: Zoom in on the dip for the coordinated control case.

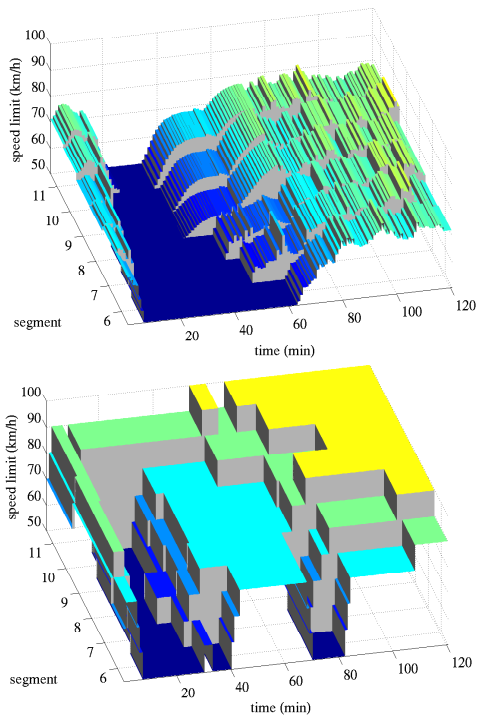


Figure 6: The speed for the continuous case without safety constraints and $N_p = 11$, $N_c = 8$ (top). The speed limits for the discrete (“ceiling”) case with safety constraints and $N_p = 11$, $N_c = 8$ (bottom).

its (base case) is effective against shock waves. The performance loss caused by discrete-valued speed limits and the inclusion of safety constraints was examined. The performance of the discrete-valued, safety constrained speed limits was comparable with the base case if the discrete-valued speed limits are generated by the “round” or “ceiling”, which are very fast compared to full integer programming. In all of the 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 a significant decrease of the total time spent.

Topics for further research include: further examination of the trade-off between efficiency and optimality for rounding versus full discrete optimization; investigation of the ef-

fectiveness of MPC for optimal coordination of speed limits for a wider range of scenarios, networks, traffic flow models and/or model parameters; and including extra control measures in addition to speed limits (such as ramp metering, dynamic lane assignment, route information, etc.).

Acknowledgments

Research supported by the Traffic Research Centre (AVV) of the Dutch Ministry of Transport, Public Works and Water Management, and by the NWO-CONNEXT project AMICI.

References

- [1] A. Alessandri, A. Di Febbraro, A. Ferrara, and E. Punta. Nonlinear optimization for freeway control using variable-speed signaling. *IEEE Trans. on Veh. Techn.*, 48(6):2042–2052, 1999.
- [2] P. Breton, A. Hegyi, B. De Schutter, and H. Hellendoorn. Shock wave elimination/reduction by optimal coordination of variable speed limits. *Proc. IEEE 5th Int. Conf. on Intelligent Transportation Syst. (ITSC’02)*, p. 225–230, Singapore, Sept. 2002.
- [3] C.C. Chien, Y. Zhang, and P.A. Ioannou. Traffic density control for automated highway systems. *Aut.*, 33(7):1273–1285, 1997.
- [4] A. Di Febbraro, T. Parisini, S. Saccone, and R. Zoppoli. Neural approximations for feedback optimal control of freeway systems. *IEEE Trans. on Veh. Techn.*, 50(1):302–312, 2001.
- [5] B.S. Kerner and H. Rehborn. Experimental features and characteristics of traffic jams. *Ph. Rev. E*, 53(2):R1297–1300, 1996.
- [6] A. Kotsialos, M. Papageorgiou, and A. Messmer. Optimal coordinated and integrated motorway network traffic control. *Proc. 14th Int. Symposium of Transp. and Traffic Theory (ISTTT)*, p. 621–644, Jerusalem, Israel, 1999.
- [7] H. Lenz, R. Sollacher, and M. Lang. Standing waves and the influence of speed limits. *Proc. Eur. Control Conf. 2001 (ECC’01)*, p. 1228–1232, Porto, Portugal, Sept. 2001.
- [8] M.J. Lighthill and G.B. Whitham. On kinematic waves: I. Flood movement in long rivers. *Proc. Royal Society of London*, 299A:281–316, May 1955.
- [9] J.M. Maciejowski. *Predictive Control with Constraints*. Prentice Hall, Harlow, England, 2002.
- [10] M. Papageorgiou. *Applications of Automatic Control Concepts to Traffic Flow Modeling and Control*. Springer, 1983.
- [11] M. Papageorgiou, J.M. Blosseville, and H. Hadj-Salem. La fluidification des rocade de l’Île de France: Un projet d’importance. Techn. Rep. 1998-17, Dynamic Systems and Simulation Lab., Techn. Univ. of Crete, Chania, Greece, 1998.
- [12] 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 I: Modelling. *Transp. Res. Part A*, 24(5):345–359, Sept. 1990.
- [13] S. Smulders. Control of freeway traffic flow by variable speed signs. *Transp. Res. Part B*, 24(2):111–132, 1990.
- [14] A.R.M. Soeterboek. *Predictive Control – A Unified Approach*. Prentice Hall, Englewood Cliffs, New Jersey, 1992.
- [15] E. van den Hoogen and S. Smulders. Control by variable speed signs: Results of the Dutch experiment. *Proc. 7th Int. Conf. on Road Traffic Monitoring and Contr.*, p. 145–149, London, 1994.