Improving Traffic Efficiency With Lane Guidance Based on Desired Speeds

Niharika Mahajan, Andreas Hegyi, Serge P. Hoogendoorn, and Bart van Arem, *Senior Member, IEEE*

***Abstract*— Drivers initiate a discretionary lane change when they perceive an anticipated improvement in their own driving condition from moving to another lane. However, such a lane change can slow down other vehicles on the target lane, and even worse initiate a disturbance. In this work, we argue that the blocking effect triggered by individual lane changes results from the heterogeneity in the desired speeds of vehicles, and thus using desired speed information of vehicles when regulating lane-changing decisions can improve traffic efficiency. In doing so, our work also exemplifies the usefulness of incorporating user preferences into control decisions. The proposed lane guidance system uses an optimization-based approach to update the target range of desired speeds on each lane in real time, and accordingly recommends individual lane changes. The control system coordinates the lane-changing decisions at the link level, for which the road stretch is subdivided into multiple sections that are controlled independently. We evaluate the performance of the lane guidance system in micro-simulation, for different network demands and desired speed distributions. The results highlight that the proposed approach utilizing the desired speed preferences of drivers results in positive efficiency gains for most traffic compositions in free flow. Moreover, the highest gains are expected in medium to high demand, and when the traffic composition includes a higher proportion of vehicles desiring higher speeds. The gains also increase when the desired speeds of vehicles that want to drive fast and those that want to drive slower are more separated.**

***Index Terms*— Lane guidance, cooperative lane-changing, desired speeds, user preferences in traffic control, traffic efficiency.**

I. Introduction

BEHAVIOUR of drivers in traffic is often self-interested. This holds also for lateral movements, wherein rational drivers try to make lane changes that maximize their own individual utility, such as keeping a desired speed, lane, or route. Lane changes are typically classified as mandatory and discretionary. While mandatory lane changes are generally coupled with more strategic objectives, like merging at a lane drop, or merging or diverging at ramps in order to keep a desired route, discretionary lane changes are less urgent and usually initiated by drivers to improve their anticipated driving condition [1]. However, the collective effect of these individual lane changes may not necessarily be optimal for the overall traffic efficiency.

Different theories have associated lateral driving behaviour to reduced freeway capacity. Laval and Daganzo [2] attribute the inefficiency to lane-changing vehicles acting like slow moving bottlenecks, creating voids ahead of them in traffic streams. Another study postulates that the effective density of the road is higher when vehicles are overtaking, because while a vehicle is executing a lane-change manoeuvre it effectively occupies its original as well as the destination lane [3]. At a macroscopic-level, some empirical studies have observed that the distribution of vehicle densities and speeds over the lanes is inefficient in saturated traffic regime [4], [5], [6]. These studies attribute the inefficiency to individual lane choice behaviour, where more vehicles prefer to stay on the faster left lanes as the total demand increases. The result is underutilised slower lanes, which operate below their individual lane capacities as the total road section nears capacity. While all these processes deteriorate traffic efficiency, we aim to regulate discretionary lane changes in a way that improves capacity by minimizing the bottleneck effect of lane-changing vehicles.

The growing ubiquity of in-vehicle systems has made it possible to both examine and to regulate lane changes at a much higher granularity. More so, lane keeping and lane change support systems [7] are becoming a regular feature in the next generation of vehicles. With these advances in mind, we develop a lane guidance system that can achieve more efficient lane use behaviour. The control approach exemplifies the value of incorporating vehicle-driver specific attributes into lane-changing decisions. Such user attributes can be directly related to the vehicle specifications, driver preferences or decisions, such as acceleration capacity of a vehicle, its route choice, or desired speed preference. These attributes can also be derived, such as driver aggressiveness, lane use behaviour, or a vehicle's level of cooperativeness in traffic. In our case, we demonstrate how individual desired speeds can be included when making lane-change decisions.

Desired speed or free speeds of a vehicle-driver combination is defined as the speed at which a vehicle drives when it is not influenced by other drivers [8]. Majority of microscopic traffic flow models assume a desired speed distribution, from which the desired speeds of individual vehicles are sampled. This is considered important for simulating realistic traffic dynamics that emerge due to the heterogeneity in individual speed preferences. Thus, in the past, some approaches to

Manuscript received 6 June 2019; revised 24 June 2020 and 21 July 2021; accepted 26 July 2023. This work was supported by Technolution B.V. under the Spookfiles A58—Phase 2 Project. The Associate Editor for this article was B. De Schutter. *(Corresponding author: Niharika Mahajan.)*

The authors are with the Department of Transport and Planning, Faculty of Civil Engineering and Geosciences, Delft University of Technology, 2628 CN Delft, The Netherlands (e-mail: niharika153@gmail.com; a.hegyi@tudelft.nl; s.p.hoogendoorn@tudelft.nl; b.vanarem@tudelft.nl).

Digital Object Identifier 10.1109/TITS.2023.3340881

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

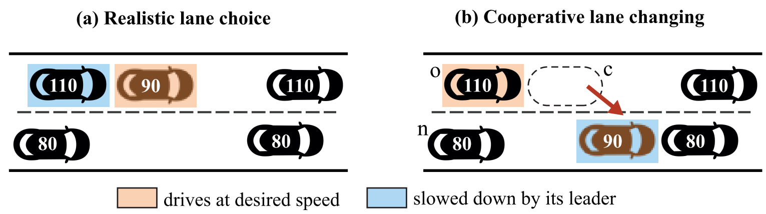

Fig. 1. An illustration of the influence of desired speeds on the impact of lane change decisions. The numbers atop the vehicles indicate their desired speeds; the figure on the right side uses labels c, o and n to identify the vehicle considering a lane change, its old follower on the original lane, and its (potential) new follower on the destination lane, respectively.

estimate desired speed distribution based on cross-sectional traffic data from loop detectors have been proposed [8], [9], [10]. Nowadays this information can be more easily acquired with in-vehicle technology, for instance, with an on-board computer or a smart phone application that determines the vehicle's desired speed based on its recent driving behaviour, not to mention cruise control systems that allow the driver to set a desired speed at which the vehicle then drives automatically.

We hypothesise that the desired speeds of vehicles driving on a multi-lane freeway have a significant influence on not just their lane choice behaviour, but also on how their lane-change manoeuvres impact surrounding traffic. Figure 1 illustrates the impact of a vehicle's desired speed and lane choice on the speed of nearby vehicles. A vehicle driving below its desired speed on a slow lane has a speed incentive to choose a faster lane. However, if the desired speed of this vehicle is lower than the average speed of the oncoming traffic on its destination lane, then a *'blocking effect'* is triggered. Since the desired speeds in traffic are heterogeneous [8], vehicles change lanes in order to maintain their desired speeds and the resulting conflicts degrade traffic performance.

The problem of lane assignment itself has been investigated in the context of automated highway systems since the 90s. The strategies developed thus far consider different traffic situations and employ different control principles to regulate lane change decisions. The earliest works [11], [12] used individual vehicle routes, more specifically information about the entry and exit ramps, to assign lanes so as to minimise path conflicts in merging and diverging traffic. The principle results in a path planning that assigns vehicles nearing their next exit to the rightmost lane available. Ramaswamy et al. [13] additionally considered the effect of traffic conditions on the choice of destination lane, within both a linear and quadratic programming problem formulation. Some other strategies develop heuristics to mitigate congestion, and associated spillback and capacity drop effects, by directing oncoming traffic away from the congested lanes [14], [15], [16]. Such heuristic approaches try to identify advice rules based on triggers in the traffic state, for situations like an activated lane drop bottleneck, over-saturated merging lane, or an incident leading to lane closure. In yet another work [17], a game-theoretic approach to improve traffic outflow by allowing vehicles to use large gaps on a right lane to accelerate before returning back to their original lanes is proposed. Although overtaking from the right is currently only permitted under congested traffic conditions (in Europe), it could potentially be legalised if intelligent vehicle systems can guarantee safety while undertaking such manoeuvres. Finally, some more recent works employ a principle based on critical lane densities. The control logic is to prevent very high lane densities by regulating traffic towards lane-specific density set-points [18], [19], [20], which may be constant values or varied according to a pre-specified function. In [21], the authors propose an extremum seeking algorithm to determine critical density set-points for such strategies.

Besides employing different control mechanisms, strategies also differ in their modelling assumptions, and solution approaches used to tackle the high dimensionality of the problem. One clear distinction is in the granularity at which lane choice decisions are modelled. Approaches that preserve the control signal as individual lane changes simplify design: by linearising the optimization problem [12], [13], [22]; by decentralising the control task among multiple link-level controllers [22], [23]; and/or using metaheuristic techniques like genetic algorithms for model or simulation-based optimization [17], [24], [25]. Alternative approaches determine a 'high-level' control signal, like lateral flows between adjacent cells in a discretised multi-lane network representation, described by a macroscopic traffic model [18], [19]. Such an aggregated control signal is typically translated to implementable control decisions using heuristics. The approach in [18] involves determining an average lane change frequency to achieve the optimal lateral flows, and then selecting the lane-changing vehicles on the basis of the gap sizes available to vehicles.

Our work contributes to this literature by developing a lane guidance system that utilises the desired speed information of vehicles in making lane-changing decisions. The control principle involves limiting the *blocking effect* triggered by self-interested lane changes that slow down other vehicles on its destination lane. At the same time, the control system may also recommend altruistic lane changes from a faster to slower lane (wherein the vehicles may not be able to keep their desired speed) when its beneficial for the total system. In practice, pursuing vehicles to realise altruistic lane changes would require some incentives, but for now we assume full compliance, in the sense that a vehicle complies to an advised lane change as long as it can find a safe gap. The focus is therefore on developing a control algorithm whereby cooperative lane changes can be made towards a system-level objective. Note that the proposed strategy is designed for under and near-saturated traffic conditions wherein a vehicle's desired speed governs its driving behaviour, unlike in over-saturated conditions where speeds are governed by the congestion dynamics. The high dimensionality of the problem is managed by using: a) multiple link-level controllers that solve an optimization problem independently to provide lane change advice to vehicles currently on their section, and b) a parsimonious lane speed model, which predicts the effective lane speed (used for estimating the objective function of the optimization problem) by using only the desired speeds of vehicles on a lane-section. The simplicity of the model ensures low computational costs, making it apt for real time control. We benchmark the performance of the lane guidance system by comparing its performance to a MOBIL-based lane change controller under different demand conditions. The choice of the MOBIL model as the basis for comparison was motivated by its formulation that uses a politeness factor to weigh the effect of a lane change decision on the utility of other drivers with its own utility. Such a MOBIL-based controller thus simulates locally cooperative lane change behaviour, and comparing our approach to it helps in assessing the value of extending the cooperation to the scale of a controlled section.

The remainder article is organised as follows. We first elaborate the functioning of the lane guidance strategy, and formulate the control problem mathematically in Section II. Subsequently in Section III, we design simulation-based experiments to evaluate the performance of the control strategy. One of the focuses here is on validating the design assumptions in our approach, and the other focus is on examining the traffic conditions under which the efficiency benefits from the optimal lane guidance system are maximised. Finally, we discuss the key findings in Sections IV and V.

II. Desired Speed Based Dynamic Lane Guidance

In this section, we propose the design of a model-based feedback control approach that results in an optimal lane allocation on a homogeneous freeway section without on-ramps, off-ramps or lane drops. Figure 2 gives an overview of the control scheme. Here we assume that individual vehicles can provide information about their desired speeds and lane positions to their link controller in real time. Using this information, the link controller determines a desired speed threshold for each lane by solving an optimization problem that maximizes a traffic efficiency criterion at each control time step and recommends lane changes to individual vehicles accordingly. The resulting traffic state is measured again in the next time step, which closes the feedback loop. The lane speed thresholds are controlled such that the slower vehicles drive on a lane more to the right (assuming right-hand driving), and the thresholds for each adjacent pair of lanes determine the range of admissible desired speeds on that lane. Vehicles with desired speeds outside the target range are advised to move to a faster or slower adjacent lane, and conversely, vehicles with desired speeds within this range are advised to keep their current lane. The mathematical formulation of the control approach is elaborated next.

A. Control Signal

The control signal in the lane guidance system are the minimum desired speed thresholds per lane, denoted as $\mathbf{u} = [u_1, u_2, \ldots, u_i, \ldots, u_{I+1}]$, where the index of each lane is

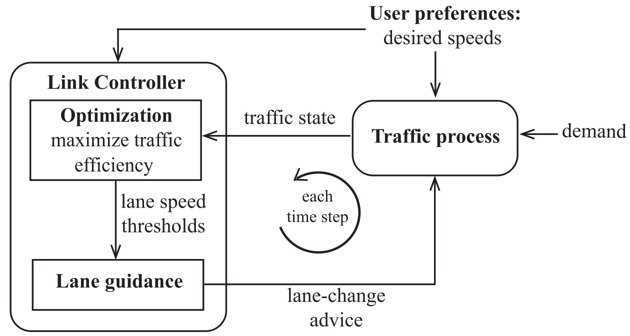

Fig. 2. Overview of the lane guidance strategy.

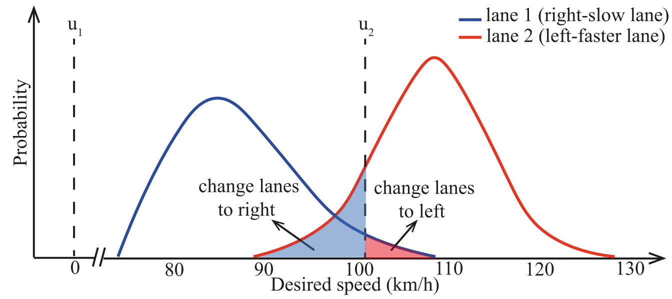

Fig. 3. The concept of desired speed thresholds for a 2-lane section with an assumed bimodal desired speed distribution. The lane change advice indicated by the blue and red shaded areas is based on the threshold values u_1 and u_2 on lane 1 and 2, respectively.

given as $i = 1, 2, \ldots, I$, such that the rightmost lane is indexed 1 and the leftmost lane I. The threshold values are increasing with the lane index: $u_1 \leq u_2 \leq \cdots \leq u_{I+1}$, and u_1 and u_{I+1} take default values $u_1 = 0$ and $u_{I+1} = \infty$, respectively. As illustrated in Figure 3, the concept of lane-specific speed thresholds requires that vehicles on the lane i with desired speed lower than u_i are instructed to changes lanes to the right, and those with desired speed higher than u_{i+1} to change lanes to the left. The only exception is if a vehicle is already executing a lane change, in which case any new advice is withheld until the ongoing manoeuvre is completed.

Furthermore, the control system is discrete-time, where the optimal control signal $\mathbf{u}^*(k)$ is evaluated by solving an optimization problem at fixed time intervals of T, chosen in the order of seconds. Here, index k refers to the time period $[kT, (k+1)T)$ when the control signal $\mathbf{u}^*(k)$ is applicable.

B. Optimization Problem

The control optimization problem is formulated to maximise total traffic flow. We use Edie's generalised definition of flow, which defines flow for a region in space and time [26]. According to the definition, flow in a region of length X and duration T depends on the distance d_n that a vehicle with index n travels and the area of the considered space-time region (refer to Figure 4) as

$$q(k) = \frac{\sum_n d_n(k)}{XT} = \frac{D(k)}{XT}. \quad (1)$$

Therefore, maximising the numerator, i.e. the total travelled distance $D(k)$ over a controlled section of length X in the

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

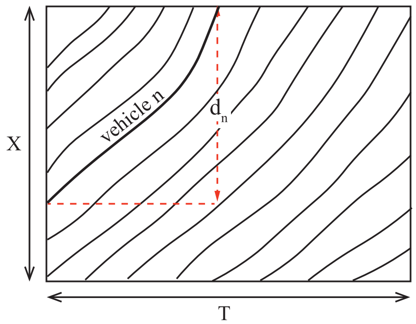

Fig. 4. Vehicle trajectories in location-time region (X-T) showing the variables in Edie's generalised definition of flow.

control time step k, would maximise flow in that section. Thus, in order to improve traffic efficiency, the optimization problem specifies the predicted total travelled distance as its objective function, which is estimated based on the current traffic state and the scontrol signal. Note that if there is no route choice possible in the controlled network, and there is no capacity drop, maximizing total flow or equivalently total travelled distance is the same as minimizing the total time spent (TTS) by vehicles in the section.

In order to determine the objective function, we predict the lane changes that would result from applying the control signal $\mathbf{u}(k)$. Traffic measurements required for this include the total number of vehicles on the I lanes, denoted as a vector $\eta = [\eta_1, \eta_2, \ldots, \eta_i, \ldots, \eta_I]$, and individual desired speeds on each lane i, denoted as a vector $\mathbf{v_i} = [v_{i,1}, v_{i,2}, \ldots, v_{i,n}, \ldots v_{i,\eta_i}]$. We can calculate the number of vehicles $f_i^R(k)$ entering lane i from its right, i.e., lane $i-1$ as

$$\begin{aligned} f_i^R(k) &= \sum_{n=1}^{\eta_{i-1}(k-1)} 1_{x_i(n)}, \\ x_i(n) &= \left\{v_{i-1,n}(k-1) \geq u_i(k)\right\}. \end{aligned} \tag{2}$$

The notation above uses an indicator function of form 1_A, where $1_A = 1$ if condition A is true, $1_A = 0$ otherwise. Similarly to (2), the number of vehicles entering lane i from the left can be estimated as

$$\begin{aligned} f_i^L(k) &= \sum_{n=1}^{\eta_{i+1}(k-1)} 1_{y_i(n)}, \\ y_i(n) &= \left\{v_{i+1,n}(k-1) < u_{i+1}(k)\right\}. \end{aligned} \tag{3}$$

With these lateral movements, we can predict the number of vehicles $\hat{\eta}_i(k)$ (the hat indicates an estimate of the quantity as opposed to an actual measurement) that would occupy lane i if the control signal $\mathbf{u}(k)$ were applied. By dropping the terms for non-existing lanes, the expression becomes

$$\hat{\eta}_i(k) = \eta_i(k-1) + f_i^R(k) + f_i^L(k) - f_{i+1}^R(k) - f_{i-1}^L(k). \tag{4}$$

In the above equation, we assume that there are enough gaps available for vehicles to safely perform all of the lane changes advised as per the control signal. In reality, some vehicles may not be able to initiate the recommended lane changes in the same time step as they receive advice in. Such vehicles would then continue to occupy their original lanes instead of the advised target lanes. Based on new traffic measurements in the subsequent time step(s), the vehicles may receive the same advice until they can perform the lane change (provided that the same advice remains the optimal one). Thus, the feedback structure in part compensates for the simplifying assumption made in Equation (4).

Further, the controller employs a parsimonious model to describe the effective lane speeds. The model describes the effective speed $V_i(k)$ of a lane as the desired speed of the slowest vehicle on it. The reasoning is that in busy traffic, the vehicle with the lowest desired speed would act as a moving bottleneck, resulting in other vehicles with higher desired speeds to drive in car-following regime. The effective lane speed $\hat{V}_i(k)$ becomes

$$\hat{V}_i(k) = \min \hat{\mathbf{v}}_\mathbf{i}(k), \tag{5}$$

where vector $\hat{\mathbf{v}}_\mathbf{i} = [\hat{v}_{i,1}, \hat{v}_{i,2}, \ldots, \hat{v}_{i,n}, \ldots, \hat{v}_{i,\hat{\eta}_i}]$ are the desired speeds of vehicles on the section assuming that the lane changes in response to the control signal $\mathbf{u}$ have been performed. Essentially, the relevant lateral movements expressed in (4) are accounted for here. The accuracy of lane speeds predicted with this simple model depends on the total number of vehicles on the section and the relative position of the slowest vehicle. We expect it to be accurate, one, when the section density is not too low, such that there are no large gaps for vehicles to accelerate without being hindered by downstream vehicles, and two, when the controlled section is neither too short that the few observed desired speeds are not representative, nor too long so as to not consider prematurely the bottleneck effect of a slow vehicle that is far downstream. We test the accuracy of this model in our simulations.

The total travelled distance $\hat{D}(\mathbf{u}(k), \mathbf{v}(k-1))$ can now be determined based on (2)-(5) as follows

$$\hat{D}(\mathbf{u}(k), \mathbf{v}(k-1)) = T \sum_{i=1}^{I} \hat{\eta}_i(k)\hat{V}_i(k), \tag{6}$$

where $\mathbf{v}(k-1) = (\mathbf{v_1}(k-1), \mathbf{v_2}(k-1), \ldots, \mathbf{v_I}(k-1))$ contains the desired speed vectors of all lanes at time step $k-1$.

Using the above elements, the optimization problem can be defined mathematically as

$$\underset{\mathbf{u}(k)}{\text{maximise}} \;\; \hat{D}(\mathbf{u}(k), \mathbf{v}(k-1)) \tag{7a}$$

$$\text{subject to} \;\; 0 \leq u_i(k) \leq V^{\text{sl}}, \quad i = 1, 2, \ldots, j \tag{7b}$$

$$u_i(k) < u_{i+1}(k), \quad i = 1, 2 \ldots, j \tag{7c}$$

$$u_i(k) = \infty, \quad i = j+1, j+2, \ldots, I \tag{7d}$$

$$u_1(k) = 0, \quad u_{I+1}(k) = \infty, \tag{7e}$$

$$\hat{\eta}_i(k) \leq l\rho_i, \quad \hat{\eta}_i(k) \in \mathbb{Z}^* = \{0\} \cup \mathbb{Z}^+. \tag{7f}$$

The problem is non-convex with a non-linear objective function and linear constraints. Constraint (7b) limits the speed thresholds $\mathbf{u}$ to be positive but below the the legal speed limit V^{sl}. Constraint (7c) ensures that the speed thresholds are increasing, from the slowest to the fastest lane. Lane index j is introduced in constraint (7d), allowing for solutions in which only lanes $1, 2, \ldots, j$ are used, and lanes $j+1, j+2, \ldots, I$ remain empty. This constraint is expected to be triggered in low demand conditions, where vehicles need not use all available lanes. Constraint (7e) specifies the boundary condition for the speed thresholds on the outermost lanes. Finally, the constraint (7f) guarantees that the estimated number of total vehicles $\hat{\eta}_i$ on a given lane section of length l does not result in a density higher than the critical lane density ρ_i. Note that the critical lane densities $\boldsymbol{\rho} = [\rho_1, \rho_2, \ldots, \rho_I]$ are traffic flow parameters that can be tuned using historic traffic data.

There may be situations where the density of a section momentarily exceeds its critical density. The last constraint is by definition violated when this happens, since the total number of vehicles in the controlled section exceeds the maximum desired occupancy

$$\sum_{i=1}^{I} \hat{\eta}_i(k) > \sum_{i=1}^{I} l\rho_i. \tag{8}$$

So, in order to minimise further disturbances induced by lane changes, the control system does not determine lane changes based on the optimization problem in (7). Instead, it recommends lane changes to distribute the surplus vehicles equally over all lanes. The underlying reasoning is that lane speeds are more uniform at higher densities (and less dependent on desired speeds of vehicles), and thus flow may be maximized by maintaining a homogeneous lane distribution while limiting the number of lane changes. The desired number of vehicles on each lane can be calculated as

$$\hat{\eta}_i(k) = \begin{cases} \left\lfloor \dfrac{\sum_{i=1}^{I} \eta_i(k-1)}{I} \right\rfloor + 1, & \text{if } i = 1, 2, \ldots j \\ \left\lfloor \dfrac{\sum_{i=1}^{I} \eta_i(k-1)}{I} \right\rfloor, & \text{if } i = j+1, j+2, \ldots, I \end{cases} \tag{9}$$

where

$$j = \sum_{i=1}^{I} \eta_i(k-1) - \left\lfloor \frac{\sum_{i=1}^{I} \eta_i(k-1)}{I} \right\rfloor I,$$

and $\lfloor\ \rfloor$ denotes a floor function.

Based on (9), the controller identifies the necessary lane changes between adjacent lanes sequentially, starting from the rightmost towards the leftmost lane. The rule is straightforward – for a deficit of n vehicles on a lane, n vehicles with the lowest desired speeds on an adjacent left lane are advised to change lanes to this lane, and for a surplus of n vehicles on the lane, n vehicles with the highest desired speeds are advised to change lanes to the adjacent left lane.

C. Solution Approach

Given the effective lane speed model in (5), the optimization problem can be solved as a partitioning problem. The task then is to partition a sorted sequence of unique desired speeds of vehicles on a controlled section into at most as many subsets as the number of lanes on the section. Note that any speed threshold between two nearest values of desired speeds will result in the same partition, and so for simplicity, we consider the candidate threshold values as halfway between consecutive (increasing) values of desired speeds reported by all vehicles on the section. For a total of N vehicles on I lanes, there are at most $N-1$ candidate threshold values, of which at most $I-1$ values are to be selected. The maximum number of possible combinations for the lane thresholds is $\sum_{k=1}^{I-1}\binom{N}{k}$. So for instance, if there were 50 vehicles on a 1 km long 3-lane controlled section, then a total of 1275 candidate solutions should be examined. Thus, given that the number of vehicles on a controlled section is not too large (which depends on the length of the section, the number of lanes, and traffic condition), the solution for the partitioning problem can be enumerated in real time. The pseudo-algorithm for solving the optimization problem by enumeration is given in Algorithm 1.

Algorithm 1 Pseudo-Algorithm for Solving the Optimization Problem by Enumeration

Input: latest measurements of desired speeds of vehicles on a controlled section, $\mathbf{v(k-1)}$
Input: number of lanes on the controlled section, I
Output: optimal lane speed thresholds to be applied in the next time step, $u^*(k)$

initialise $maxTTD = 0$;
$uniqueDesSpeeds \leftarrow$ unique and sorted (in ascending order) desired speeds values in $\mathbf{v(k-1)}$;
$candidateThresholds \leftarrow$ average of consecutive elements in $uniqueDesSpeeds$;
$thresholdCombinations \leftarrow$ set of all possible combinations for choosing $\{1, 2, \ldots, I-1\}$ elements from $candidateThresholds$;
if total number of vehicles on the section is below a maximum value (checked in (8)) **then**
  foreach $combination$ in $thresholdCombinations$ **do**
    $TTD \leftarrow$ total travelled distance as per (6) for lane thresholds in $combination$;
    if $TTD > maxTTD$ **then**
      $maxTTD \leftarrow TTD$;
      $optCombination \leftarrow combination$
    end
  end
  $u^*(k) \leftarrow$ padded vector $[0, optCombination, \infty \overset{\times N}{\cdots\cdots}]$, where $N = I - dim(optCombination)$;
else
  advised lane changes ensure that the number of vehicles on each lane are as per (9)
end

III. Simulation-Based Evaluation

In this section, we evaluate the lane guidance strategy. To begin with, we made some simplifying modelling assumptions to predict the objective function of the control-optimization problem. So, first we test the accuracy of the predictions against simulated data. Next, we validate the control response that leads to an improvement in traffic performance. Moreover, our approach has similarities to MOBIL - an accepted lane-changing model [1]. Thus, we benchmark the performance of our controller against a MOBIL-based controller under different demand conditions. Next, we compare the controller gains for different desired speed scenarios, since the distribution of desired speeds in traffic is not well established empirically and can vary depending on the traffic context. Finally, we assess the impact of having an additional lane in the network on the control performance.

The lane guidance system is implemented in the VISSIM microsimulation tool. This means that the designed controller overrides the desired lane changes determined by the default lane change model in VISSIM. However, the execution of a lane change decision is controlled by VISSIM, meaning that VISSIM decides to initiate a recommended lane change when a gap becomes available on the destination lane. In the base case, however, the lateral manoeuvres are handled entirely by VISSIM based on its empirical-psychophysical lane change model. The vehicles furthermore behave in accordance with the 'slow lane rule' - this driving behaviour type follows the German Traffic Code (representative of the European overtaking rules), mandating a vehicle to overtake using a faster (left) lane, and generally keeping to the right as much as possible. The car-following behaviour uses the Wiedemann 99 model with the default parameter values specified in VISSIM.

The strategy is applied over a 5 km long 2-lane freeway in all experiments, except ones in Section III-E, where a similar 3-lane freeway is used. The following parameter setting is used in the control implementation

- $\rho_1 = 35\ veh/km$ and $\rho_2 = \rho_3 = 30\ veh/km$; critical densities on lanes 1 (slowest), 2 and 3, respectively
- $l = 1\ km$; length of each controlled section
- $T = 5\ s$; control time-step size

The critical lane density values are tuned manually based on fundamental diagrams plotted with simulation data.

We model the heterogeneity in desired speeds of vehicles by specifying 3 different vehicle speed classes. These classes represent the diversity in mixed traffic. Class 1 includes fast passenger or commercial cars with desired speeds closer to the legal speed limits (for instance 130 km/h in the Netherlands); class 2 includes cars that desire relatively lower speeds due to driver or vehicle characteristics; and class 3 represents trucks or heavy vehicles whose speeds are legally limited to lower values (for instance to 80 km/h in the Netherlands). We consider 4 scenarios employing different distributions for the 3 speed classes, as illustrated in Figure 5. Scenario 1 is the most straightforward, using discrete probability distributions for all classes. Differences in the desired speeds between the classes result in conflicts as vehicles try to maintain their own desired speeds. We use this scenario to validate the control behaviour, which we expect should minimize such conflicts. Scenarios 2 and 3 are designed such that in one the desired speeds of the car classes are set apart more than the other. Finally, Scenarios 4 and 5 employ truncated normal distributions to test the effect of a non-uniform distribution on the control performance.

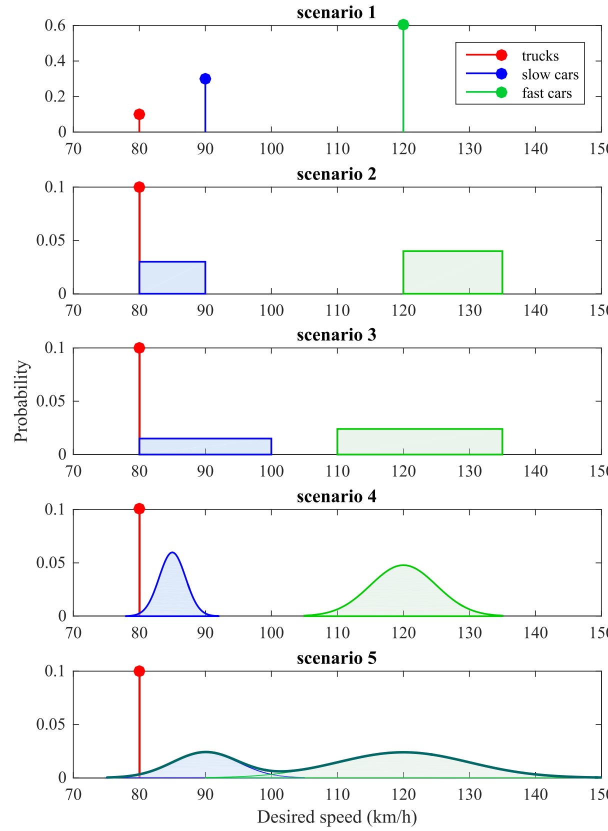

Fig. 5. Scenarios with different desired speed distributions for each vehicle speed class. The traffic composition in all scenarios is: [60%, 30%, 10%] of [fast cars, slow cars, trucks]. For scenario 5, the composite desired speed distribution is shown by the dark green line.

A. Accuracy of Total Travelled Distance Predictions

The optimization problem defined in Section II-B minimises the total travelled distance D in the controlled section. The estimation of D is based on the parsimonious lane speed model, which makes solving the optimization problem computationally suitable for a real-time implementation. However, the prediction accuracy of a simplistic prediction model is unknown. We evaluate the model accuracy by comparing the prediction estimates to the corresponding real values as obtained from simulated vehicle trajectories. In this evaluation, we used scenario 4 as shown in Figure 5, and a staircase demand function with flow values ranging between 1500 and 4000 veh/h, with a flow step of 300 veh/h every 300 s for a total of 3300 s. The lane change control is active for half of this duration, so the prediction accuracy reflects both with and without control situations.

The results in Figure 6 show how the model performs; each data-point in the plots are the estimate for the total travelled distance by vehicles in a controlled section in a single sampling time step. We used a sampling time of 5 s for our results. Notice in the top sub-plot that the prediction estimates in general follow the actual measurements based on VISSIM data; the $R^2 = 0.92$ in the simulated scenario. However, the density of data-points under the 100% accuracy reference line is more than that above it, indicating that the model on an average tends to underestimate the total travelled distance. The finding is plausible since the assumption of the average

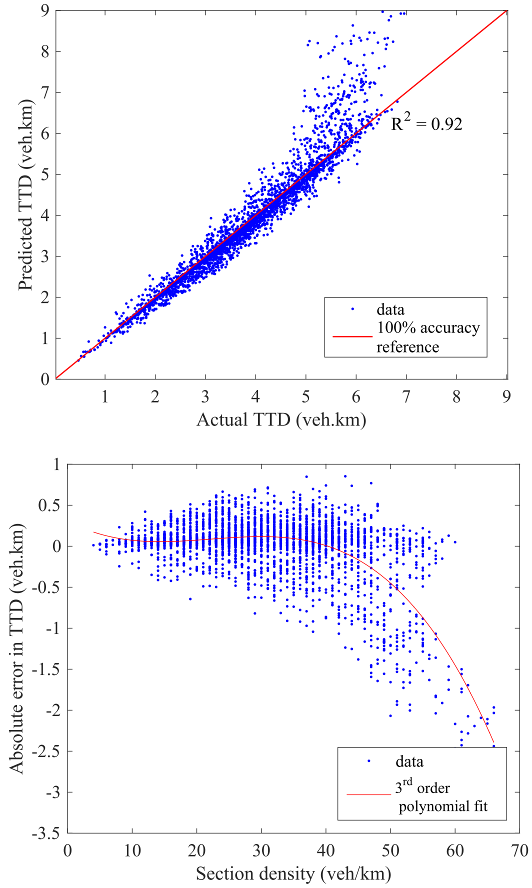

Fig. 6. Accuracy of total travelled distance predictions using the parsimonious lane speed model.

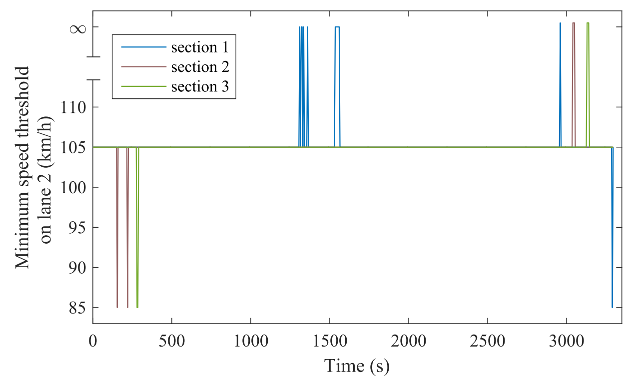

Fig. 7. Control signal showing the speed threshold on the left lane for 3 different controlled sections. Note that all sections keep a 105 km/h threshold for most of the simulated period, not visible because of overlapping plot-lines.

lane speed to be the desired speed of the slowest vehicle on the lane section is rather conservative for free flowing traffic. However, the same assumption in congested traffic results in a significant overestimation of the total travelled distance, as can be observed in bottom sub-plot. In this sub-plot, it is clear that the errors grow for densities above 50 veh/km over the 2 lanes, and the total travelled distance model is not as accurate for predictions in congested regime. To sum up, the proposed prediction model is reliable in under or near-saturated densities, which is the intended traffic regime for application of the control approach.

B. Control Behaviour

The control strategy requires each controlled section to solve the optimization problem at fixed time intervals. In this section, we investigate the control response and traffic behaviour thus achieved. The presented results are simulated for desired speed scenario 1 in Figure 5, and a staircase demand function with flows ranging between 1500 and 4000 veh/h. This results in near-capacity flows for part of the simulation duration, while ensuring that the traffic flow is never congested. Moreover, the advantage of using a scenario with discrete desired speeds distributions is that the optimal solution for lane assignment has 3 possible configurations, which are easy to interpret. The optimal speed thresholds can result in a lane distribution where: a) the slow and fast cars drive on the left fast lane while trucks drive on the right, b) all slow cars drive with the trucks on the slow-right lane, or c) all vehicles occupy the right lane. Figure 7 shows the control signal for 3 different controlled sections. Notice that the threshold values for the left-fast lane are 85, 105 km/h, and ∞, each corresponding to the 3 configurations respectively. Since the threshold value of 105 km/h prevails for most of the time, we interpret that the strategy of advising the slow cars on the left lane to change lanes to the right lane is optimal. This solution is in accordance with our expectation, because if vehicles with desired speed of 90 km/h drive on the slower right lane their speed degradation from driving behind a 80 km/h truck is less compared to the speed loss experienced by 120 km/h vehicles if they were to drive on the left lane. Note that the spikes in the control signal at 85 km/h and ∞ are related to unique situations, where the traffic composition on the section either includes no fast cars or only fast cars. In the case with no fast cars, the controller separates the slow vehicles from the trucks onto the two lanes, and when there are only fast cars on a section, they can drive on the right lane without any hindrance.

Note that when the desired speeds are continuously distributed, like in scenarios 2-4 in Figure 5, the possible configurations for distributing vehicles over the lanes based on their desired speeds increase and the optimized control signal would thus show more variability.

In order to understand how the control strategy affects traffic state, in Figure 8 we plot fundamental diagrams – showing the flow and density relationship per lane – for with and without control cases. Here it becomes clear that in the base case both lanes operate at similar speeds, approximately 85-90 km/h, whereas the lane speed difference is more significant when the optimal strategy is active. In the controlled case, the average speed on the left lane is about 120 km/h and the right lane about 80 km/h. Thus, the speed gain for the left lane is higher than the loss for the right lane, which explains the overall improvement in efficiency from lane guidance. We measure improvements in traffic efficiency in terms of percentage savings in total time spent (TTS) for the same demand pattern entering the network. For the scenario in Figure 8, our strategy results in a 7.5% improvement in TTS.

Furthermore, we investigate the lane-changing behaviour resulting from the optimal strategy. Since realisation of a lane change recommendation depends on gap availability on the

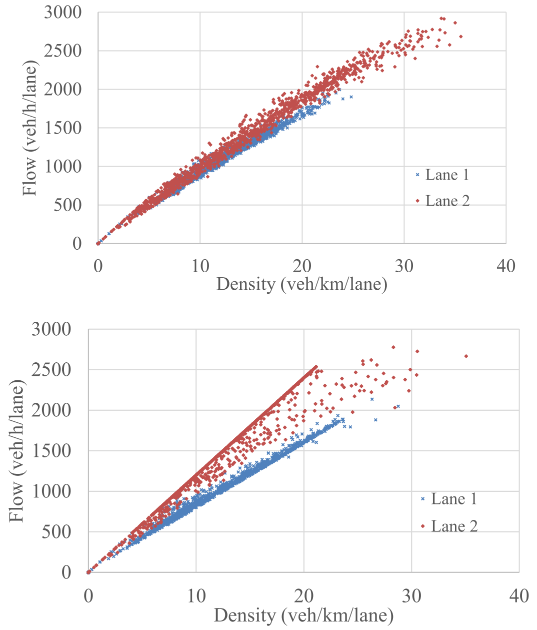

Fig. 8. Fundamental diagrams comparing traffic behaviour with VISSIM's lane change behaviour (top), and optimal lane guidance (bottom).

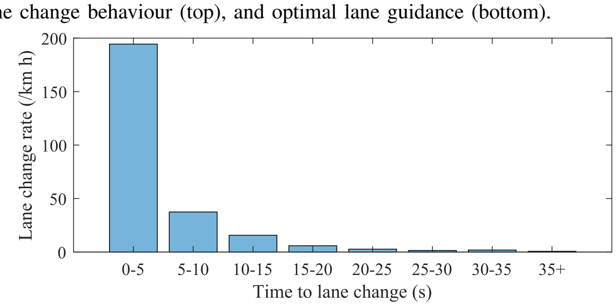

Fig. 9. Histogram plot for the time duration taken by vehicles to initiate a lane change manoeuvre after receiving an advice. The data includes 99% of all recommended lane changes, discarding the 1% that were not realised.

destination lane, the time a vehicle takes before initiating an advised lane change manoeuvre varies with the local traffic state. Figure 9 plots the frequency of different latencies in lane-changing. The results show that even though some advised lane changes were realised in over 30 s, 75% of all the lane changes were initiated within 5 s of receiving an advice. This insight validates the choice of $T = 5\ s$, as most recommended lane changes are expected to be realised before the control advice is updated in a subsequent time-step.

C. Benchmarking Control Performance

The performance of the proposed optimal lane guidance approach is compared to a microscopic controller based on the lane-changing model MOBIL: Minimizing Overall Braking Induced by Lane changes [1]. The basic principle of MOBIL is to compare the utility of the current lane with that of a prospective destination lane. These utilities are calculated in terms of accelerations, available from an underlying car-following model. In addition to its compact formulation, the model incorporates a *politeness parameter* p to capture the cooperativeness of the lane-changing vehicle towards its neighbouring vehicles, namely its current follower and the prospective new follower on its destination lane. The value of the politeness parameter can thus be varied to demonstrate different levels of cooperation; $p = 0$ would imply selfish vehicles, and $p = 1$ fair vehicles that consider the utility of other vehicles equivalent to their own.

Figure 1 labels the vehicles involved in a lane change decision – the vehicle c considering a lane change, its old follower o on the current lane, and the prospective new follower n on the destination lane. Furthermore, acceleration a_x denotes the acceleration of a vehicle x in the current lane just before initiating a lane change. The anticipated acceleration of this vehicle on the destination lane is given as $\tilde{a}_x$; we compute such future accelerations using the IDM car-following model [27]. The lane change decisions in MOBIL are evaluated using the following two criteria:

a) a *safety criterion* that ensures that the deceleration of the new follower on the destination lane is not larger than a threshold b_{safe},

$$\tilde{a}_n \geq -b_{\text{safe}}. \tag{10}$$

b) an *incentive criterion* that balances the utility gain for the lane-changing vehicle with the (dis)utility of its neighbours. The utility of the neighbouring vehicles is weighted by the politeness parameter, and a lane change is executed if the total weighted utility is larger than a threshold value Δa_{th}:

$$\underbrace{\tilde{a}_c - a_c}_{\text{lane-changer}} + p \left(\underbrace{\tilde{a}_o - a_o}_{\text{old-follower}} + \underbrace{\tilde{a}_n - a_n}_{\text{new-follower}} \right) > \Delta a_{th}. \tag{11}$$

The above criterion is adapted for asymmetric passing rules, as applicable to European driving. Firstly, in order to prevent overtaking from the right (unless traffic is congested), the utility of a right lane is bounded by the utility of the left lane in free-flow. This means that the prospective acceleration on the right lane $\tilde{a}_c^{\text{eur}} = \min(a_c, \tilde{a}_c)$, is restricted by the acceleration a_c on a current (left) lane. Secondly, in order to enforce right-hand driving wherein vehicles keep to the right lane unless overtaking, a constant bias $\Delta a_{\text{bias}} > 0$ is introduced to favour the right lanes. For the same reason, the utility of the follower on the right lane is also neglected. The result is that the vehicles by default prefer a right lane. With these adaptations to (11), a vehicle changes lanes to right if

$$\tilde{a}_c^{\text{eur}} - a_c + p\left(\tilde{a}_o - a_o\right) > \Delta a_{th} - \Delta a_{\text{bias}}. \tag{12}$$

A vehicle changes lane to the left if

$$\tilde{a}_c - a_c^{\text{eur}} + p\left(\tilde{a}_n - a_n\right) > \Delta a_{th} + \Delta a_{\text{bias}}. \tag{13}$$

MOBIL can produce locally cooperative lane change behaviour when employing $p > 0$. While local cooperation is not necessarily a given in real traffic, it is plausible since vehicles can (to some extent) anticipate the effects of their lane change on nearby vehicles. Thus, to get insight into the value of more widespread cooperation in lane-changing, we compare the performance of our optimal strategy to a cooperative MOBIL-based lane change controller. For the latter, we employ a politeness parameter value of $p = 1.0$. Similar to the implementation of our controller, here the rules of MOBIL model govern the lane-change decisions of

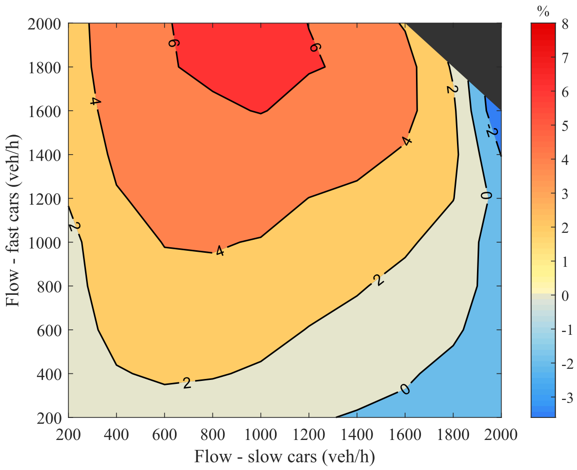

(a) Optimal Lane Guidance

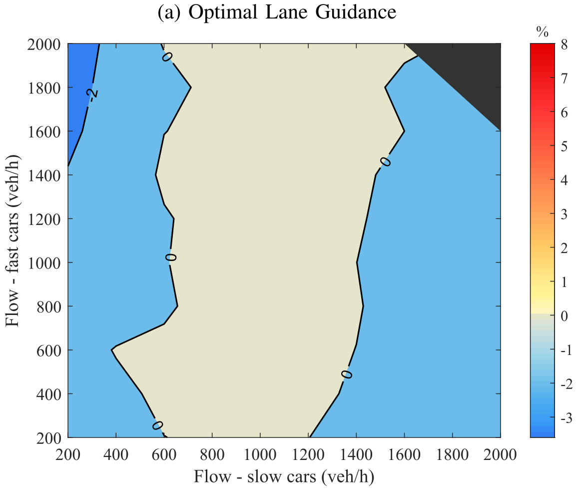

(b) Cooperative MOBIL (p = 1.0)

Fig. 10. Percentage gain in total time spent as compared to VISSIM for different traffic flow compositions (averaged results for 10 seeds). In black are scenarios in which complete demand cannot be loaded onto the network.

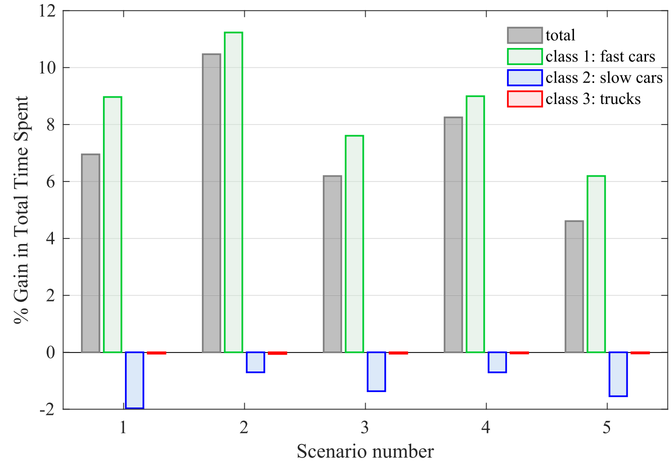

Fig. 11. Distribution of total time spent gains between vehicle classes (averaged results for 10 seeds).

individual vehicles, but the execution of the manoeuvre and longitudinal dynamics are controlled by VISSIM. We compare the performance of the two controllers for a range of traffic demands, varying flows of fast cars (class 1) and slow cars (class 2). The experiments use a constant demand, and a discrete desired speed distribution similar to scenario 1 in Figure 5.

Figure 10 shows the contour plots for the TTS improvements in the different demand scenarios. Observe that we exclude high-demand scenarios in which VISSIM cannot load the complete demand onto the network. Results for the feasible demand scenarios show that the TTS gain vary between -2 to 7% for the optimal lane guidance strategy. Even though the gains are largely positive, the strategy does have a feasible application region where the TTS savings are more significant. It is especially beneficial to use the strategy when the demand for the faster vehicle class is reasonably high (above 1200 veh/h in our experiments), while the demand for a slower vehicle class is moderate (between 400 and 1600 veh/h in our experiments). This finding matches intuition; the control approach in principle tries to balance the trade-off between the speed loss for the slow vehicles when assigned to the right lane along with (even slower) trucks, and the speed gain for the fast cars on the left lane by preventing the blocking effect of the slow cars. As the volume of the fast cars increases, the performance gain for the fast cars increasingly dominates the loss for the slow cars from driving on the slower lane.

In comparison to the optimal strategy, cooperative MOBIL strategy does not result in similar gains. An argumentation for this finding is that the cooperation is limited to one other vehicle in the European adaptation of the model (as given in Equations (13) and (12)). We find its locally cooperative behaviour to be comparable to the psychophysical lane-change model of VISSIM. Thus, the extension of cooperation among vehicles in a controlled section in our lane guidance system attributes to the efficiency gains achieved.

D. Impact of Desired Speed Heterogeneity

The performance of the lane guidance system is compared for the different desired speed distribution scenarios provided in Figure 5. For this we simulated each scenario with 10 different random seeds, using a demand profile with flows varying between 1500 and 4000 veh/h, and a traffic composition with 60% fast cars, 30% slow cars and 10% trucks. In Table I we compare the 5 scenarios based on the percentage gains in TTS, as well as characteristics of the lane change behaviour. In Figure 11 we break down the total TTS gains into gains per vehicle class, resulting in some important insights.

There is an overall TTS gain, ranging between 4-10%, in all scenarios. This gain is the result of fast cars (class 1) being able to travel faster, while the slow cars and trucks (classes 2 and 3) on average go slower and have a negative TTS gain. In scenario 2, the percentage TTS gain for fast cars is the highest and the TTS loss for slow cars is the lowest compared to other scenarios, resulting in a total TTS gain of 10.5%. Such sizeable gains can be explained by the high inter-class difference in desired speeds in this scenario, which means that the slow cars significantly impede the fast cars in the uncontrolled case. The relatively small TTS loss for the slow cars in scenario 2 is because of the low intra-class variation in desired speeds, which are not significantly different from the slowest moving trucks. In contrast, slow cars incur the highest loss in TTS in scenario 1, because all the slow cars driving

TABLE I
CONTROL PERFORMANCE FOR THE SCENARIOS IN FIGURE 5 (AVERAGED RESULTS FOR 10 SEEDS)

Scenario \ Indicator	TTS gain (in %)	No. of lane changes (in $/km\,h$)	Avg. time to lane change (in s)	Realisation* rate (in %)
1	7.0	285	3.8	99
2	10.5	318	4.9	97
3	6.2	480	3.7	92
4	8.3	439	4.3	93
5	4.6	676	3.3	89
Avg. margin of error**	± 0.5	± 19	± 0.2	± 0.6

* Realisation rate is subject to gap availability; a vehicle adheres to an advice when it can initiate a lane change within the duration that it receives the same destination lane advice continuously.

** Margin of error for the 95% confidence interval is calculated as $\frac{\rho * t_{0.975}}{\sqrt{n}}$, where ρ is the standard deviation of the indicator value in n simulation runs ($n = 10$), and $t_{0.975} = 2.262$ is the two-tailed t-score.

TABLE II
CONTROL PERFORMANCE IN A 3-LANE NETWORK (AVERAGED RESULTS FOR 10 SEEDS)

Scenario \ Indicator	Demand (in veh/h)	TTS gain$_1$* (in %)	TTS gain$_2$** (in %)	No. of lane changes (in $/km\,h$)	Avg. time to lane change (in s)	Realisation rate (in %)
2	1500-4000	7.2	14.9	1808	1.9	96
2	2500-5000	9.9	-	2189	2.3	95
5	1500-4000	2.5	10.2	2398	1.6	95
5	2500-5000	4.5	-	3008	2.0	93

* Gain in total time spent relative to baseline 3-lane performance in VISSIM.
** Gain in total time spent relative to baseline 2-lane performance in VISSIM.

behind a truck on the slower lane must drive at least 10 km/h below their desired speeds.

In scenarios 3 and 4, there is more intra-class heterogeneity in desired speeds, which means that despite control, there will be more variation of desired speeds within each lane. For wider desired speed distributions, the eventual distribution within a lane will also in general be wider. Thus, the faster vehicles will be slowed down more, depending on the desired speeds of the individual vehicles driving in a car-following regime. We see that this results in lower TTS gains for the fast cars in these scenarios compared to the first two scenarios. Another explanation is the lane change characteristics shown in Table I, which suggest that the increase in desired speed heterogeneity in these scenarios is associated with a roughly 35% increase in the lane change rate, and a slightly lower realisation rate for the advised lane changes. This effect is further aggravated in scenario 5, where the desired speed distribution for the slow and fast vehicles spread over an even wider speed range. Even so, the average time to execute an advised lane change is quite similar across all the scenarios, ranging between 3-5 s.

E. Performance in a 3-Lane Network

In this section, we analyse the effect of adding an additional lane to the network on control performance. For this we selected scenarios with the highest and lowest TTS gains in the 2-lane network, namely scenarios 2 and 5, and simulated them with two different demand profiles. The first demand profile is kept the same as that for the 2-lane simulations in Section III-D, and in the second the input flows are increased to vary between 2500 and 5000 veh/h. The results in Table II show that similar to the 2-lane network TTS gains for scenario 2 are higher compared to scenario 5. Additionally, for demand varying between 1500 and 4000 veh/h, the overall TTS gains are positive yet lower in comparison to the corresponding 2-lane cases. The reason is that the 2-lane demand represents an under-saturated traffic situation in the 3-lane network, and results in an improved uncontrolled performance based on the VISSIM lane change model. This is apparent from the significantly higher TTS gain$_2$ results, where we compare the TTS in the controlled 3-lane case to the baseline TTS in uncontrolled 2-lane network. These results suggest that the additional lane benefits the control performance notably. An extra lane implies an extra desired speed bin for the controller, and more homogeneous desired speeds per lane. More importantly, the TTS gains in the 3-lane network are higher in more saturated traffic condition for both scenarios. Even with the more heterogeneous desired speed distribution in scenario 5, there is a 4.5% gain in TTS from lane guidance.

IV. DISCUSSION

Our work proposes a desired speed-based lane guidance system that advises lane changes so as to optimize overall traffic efficiency. The incorporation of individual desired speeds is a notable advantage of the proposed control approach. To actualize this, the control system relies on vehicle-to-infrastructure communication. Vehicles provide their desired speeds and lane positions to the controller managing the link they occupy, and in return receive lane change advice. The responsiveness of the control strategy to individual speed preferences also adds to its future-readiness. As automated vehicles of varying levels of automation are adopted, we expect a transition phase where manually driven and automated vehicles would share the road infrastructure. Such mixed traffic may further increase heterogeneity in desired speeds of vehicles with different technical capabilities, making it even more important to include such user preferences in control decisions.

Even though the desired speed distribution is an important input to most micro-simulators, measuring it from real traffic data is challenging, as vehicles are not always free driving, rather are likely to be impeded by a slower leading vehicle. A few known studies have attempted to empirically estimate this distribution [8], [10], [28]. These studies typically distinguish sample data for different vehicle-types: generally passenger cars and heavy vehicles, fitting a unique distribution to each. In our simulation scenarios (refer to Figure 5), we specify independent desired speed distributions for three vehicle classes, based on speed instead of vehicle-type, although there is correlation between the two. These include two passenger car classes, namely fast cars and slow cars, in addition to trucks or heavy vehicles. The distinction of slow and fast cars was made to be able to set up controlled experiments, where the proportion of vehicles that can potentially have impeding effects on the speed of other vehicles could be modulated. Nevertheless, the proposed control approach can be applied for any (composite) desired speed distribution in traffic.

Furthermore, we think that the reported efficiency gains with the lane guidance system might be conservative estimates for possible gains in a real-world implementation. This is because of how VISSIM assigns a lane to vehicles entering the network. A vehicle gets assigned to the lane that minimises the time-to-collision to the prospective leading vehicle. As a result, we have an initial lane assignment that does not explicitly depend on the individual speed preferences, and in that sense is somewhat arbitrary. In real traffic, the lane choice of vehicles is informed by their current and desired speeds, and by the ‘keep to right/left’ driving rule. Thus, enforcing the lane speed thresholds may require fewer lane changes in practice than in our simulations. A lower lane change rate implies fewer lane change triggered disturbances, leading to better traffic stability and efficiency.

Finally, while our approach is designed to regulate discretionary lane changes over a homogeneous length of freeway, there are other approaches that specifically regulate infrastructure-related lane changes at a lane-drop, on-ramp or off-ramp. From a practical viewpoint, it is important to integrate lane-change advice from different sources to ensure control consistency and effectiveness. For instance, establishing a hierarchy of lane-change motivations can be one way to resolve conflicts. Furthermore, for an efficient lane assignment, lane change decisions that have precedence over discretionary lane changes, should also be considered in the optimization problem formulated in our approach. Another aspect is the application of the lane guidance system in over-saturated traffic condition. The control requirement in congested state is to reduce the flow towards congestion (to resolve the queue faster), which conflicts with the optimization objective to maximize total flow. Thus, once a breakdown occurs, the lane guidance system should not be deployed stand-alone. Studies have shown that integrating lane guidance with dynamic speed control can be effective in mitigating congestion [14], [16].

V. Conclusion and Future Work

Individual lane-changing decisions are important determinants of overall traffic efficiency. Research focusing on regulating lane-changing behaviour towards a system objective have seldom included vehicle-level preferences into the control decisions. In this article, we have emphasised that heterogeneity of individual desired speeds can cause inefficiency in busy traffic – where drivers prefer the faster left lane(s) and try to avoid the right lane(s). The proposed control strategy uses this information to advise lane changes that minimize the *blocking effect*, i.e. the speed reduction caused due to desired speed differences between vehicles following one another. For this we use the notion of minimum desired speed thresholds on individual lanes as the control signal, which translate to a target range for desired speeds on every lane. The controller then recommends individual lane-changes in order to assign vehicles as per the admissible desired speeds on the lanes.

Our approach used a parsimonious model that incorporated desired speed information of vehicles to predict effective lane speeds. The simplicity of the model allowed to update the control signal in real time. We validated the accuracy of the lane speed model for predicting the totalled travelled distance. Simulation results showed an overall good model fit ($R^2 = 0.92$ in the simulated scenario) to VISSIM trajectory data. It was in accordance with our expectation that the model was found to have a high accuracy in unsaturated and near-saturated traffic conditions, and deteriorated performance in congested regime – specifically for densities above 50 km/veh on the two lanes. The prediction errors were moreover found to increase as a function of density and fitted to a third-order function. Such curve-fitting done with a representative dataset, including more diverse traffic scenarios, could be used to specify a correction function for the prediction model. Note that the control approach is modular, and the proposed lane speed model could be replaced with another model that incorporates desired speed information of individual vehicles.

The simulation results have highlighted some important characteristics of the control approach. On comparing the fundamental diagrams for cases with and without control, we have shown that the underlying mechanism for improvement in traffic efficiency involves improving the speed on the faster lanes. This was also our expectation, as we try to eliminate selfish lane changes that slow down other upstream vehicles with relatively higher desired speeds. Further, we found that over 90% of the advised lane changes were realized in all simulated cases. Thus, recommending lanes to individual vehicles, without explicitly controlling the manoeuvre at the operational level, did not impede compliance to the lane guidance.

We also investigated the influence of the traffic composition on control performance in a 2-lane network, measured in terms of improvement in total time spent as compared to the VISSIM microsimulation model. Different scenarios were tested by varying the demand and desired speed distribution for the car-vehicle classes, while trucks comprised 10% of the total demand. The highest gains occurred in high demand conditions where the proportion of fast cars was also high. The strategy improved efficiency for all scenarios except when the demand for slow cars was too high (>1600 veh/h on 2 lanes). However, such a scenario where the slow cars dominate the traffic composition is not very realistic, especially on highways. Moreover, the gains depended strongly on the heterogeneity of desired speeds within and in between the vehicle speed classes. The lane guidance resulted in higher gains when the inter-class heterogeneity in desired speed was high. Conversely, increasing intra-vehicle class heterogeneity led to lower inter-class heterogeneity, resulting in relatively lower gains. Nevertheless, positive TTS gains were noted even in scenarios with the more heterogeneous desired speed distributions. Some additional experiments examined the impact of an additional lane on control performance. For lower flows, the TTS gains from lane guidance are diminished, however, with higher flows the TTS gains in the 3-lane network were comparable to gains in the 2-lane network in the tested scenarios.

In order to assess the value of cooperative lane-changing, we also compared our approach to a locally cooperative lane guidance system. The latter was based on the MOBIL model [1], in which lane-changing decisions balance the utility of lane change to the ego vehicle with the potential disbenefits it causes to the nearby vehicles. The results highlighted that the MOBIL-based controller performs very similarly to VISSIM's lane change model, and our approach significantly outperforms the MOBIL-based controller. We conclude that the extension of cooperation in lane-changing to the link level is beneficial to the traffic system.

Nonetheless, the proposed approach can be further improved. An important aspect is to improve the validity of the lane speed model for a wider range of traffic conditions, for instance by incorporating the order, position and current speed of the individual vehicles. However, such model development should still ensure sufficient computational efficiency for real-time control. Another direction for improving controller performance is to combine some gap acceptance criterion with the lane change decisions. This would imply a full realization rate, and the option to accurately include the control response (subject to gap availability) into the optimization problem. Finally, adding such complexities to the model would also need more efficient methods to solve the optimization problem.

Lastly, upon investigating the distribution of efficiency gains amongst the different vehicle classes, we found that the gains are disproportionately biased towards the faster vehicle class. The slower vehicles on an average experience higher travel time. In other words, the overall efficiency gain comes at the cost of an efficiency loss – even though small – for the slower traffic. This disparity between the vehicles based on their speed preferences must be compensated to ensure fairness. Integrating the proposed control approach with economic incentives, which guarantee both equity and wilful participation in the system, will be the topic for our future research.

References

[1] A. Kesting, M. Treiber, and D. Helbing, "General lane-changing model MOBIL for car-following models," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 1999, no. 1, pp. 86–94, Jan. 2007.

[2] J. A. Laval and C. F. Daganzo, "Lane-changing in traffic streams," *Transp. Res. B, Methodol.*, vol. 40, no. 3, pp. 251–264, Mar. 2006.

[3] W.-L. Jin, "A kinematic wave theory of lane-changing traffic flow," *Transp. Res. B, Methodol.*, vol. 44, nos. 8–9, pp. 1001–1021, Sep. 2010.

[4] S. Ahn and M. J. Cassidy, "Freeway traffic oscillations and vehicle lane-change maneuvers," in *Proc. 17th Int. Symp. Transp. Traffic Theory (ISTTT17)*. London, U.K.: The Institution of Mechanical Engineers, Jul. 2007, pp. 691–710.

[5] A. Duret, S. Ahn, and C. Buisson, "Lane flow distribution on a three-lane freeway: General features and the effects of traffic controls," *Transp. Res. C, Emerg. Technol.*, vol. 24, pp. 157–167, Oct. 2012.

[6] V. L. Knoop, A. Duret, C. Buisson, and B. van Arem, "Lane distribution of traffic near merging zones influence of variable speed limits," in *Proc. 13th Int. IEEE Conf. Intell. Transp. Syst.*, Sep. 2010, pp. 485–490.

[7] M. Tideman, M. C. van der Voort, B. van Arem, and F. Tillema, "A review of lateral driver support systems," in *Proc. IEEE Intell. Transp. Syst. Conf.*, Sep. 2007, pp. 992–999.

[8] S. P. Hoogendoorn, "Unified approach to estimating free speed distributions," *Transp. Res. B, Methodol.*, vol. 39, no. 8, pp. 709–727, Sep. 2005.

[9] H. Botma, "The free speed distribution of drivers: Estimation approach," in *Proc. Five Years Crossroads Theory Practice, 5th TRAIL Annu. Congr., 2*, Delft Univ. Press, TRAIL Research School, Delft, The Netherlands, 1999, pp. 1–22.

[10] J. Geistefeldt, "Empirical estimation of desired speed distributions for microscopic traffic simulation," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2490, no. 1, pp. 41–46, Jan. 2015.

[11] B. S. Y. Rao and P. Varaiya, "Roadside intelligence for flow control in an intelligent vehicle and highway system," *Transp. Res. C, Emerg. Technol.*, vol. 2, no. 1, pp. 49–72, Mar. 1994.

[12] R. W. Hall and C. Caliskan, "Design and evaluation of an automated highway system with optimized lane assignment," *Transp. Res. C, Emerg. Technol.*, vol. 7, no. 1, pp. 1–15, Feb. 1999.

[13] D. Ramaswamy, J. V. Medanic, W. R. Perkins, and R. F. Benekohal, "Lane assignment on automated highway systems," *IEEE Trans. Veh. Technol.*, vol. 46, no. 3, pp. 755–769, 1997.

[14] W. J. Schakel and B. van Arem, "Improving traffic flow efficiency by in-car advice on lane, speed, and headway," *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 4, pp. 1597–1606, Aug. 2014.

[15] R. V. Rabsatt and M. Gerla, "Multi-lane cooperative shock wave mitigation," in *Proc. 16th Annu. Medit. Ad Hoc Netw. Workshop (Med-Hoc-Net)*, Jun. 2017, pp. 1–7.

[16] Y. Zhang and P. A. Ioannou, "Combined variable speed limit and lane change control for truck-dominant highway segment," in *Proc. IEEE 18th Int. Conf. Intell. Transp. Syst.*, Sep. 2015, pp. 1163–1168.

[17] S. Yao, V. L. Knoop, and B. van Arem, "Optimizing traffic flow efficiency by controlling lane changes: Collective, group, and user optima," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2622, no. 1, pp. 96–104, Jan. 2017.

[18] C. Roncoli, I. Papamichail, and M. Papageorgiou, "Hierarchical model predictive control for multi-lane motorways in presence of vehicle automation and communication systems," *Transp. Res. C, Emerg. Technol.*, vol. 62, pp. 117–132, Jan. 2016.

[19] C. Roncoli, N. Bekiaris-Liberis, and M. Papageorgiou, "Lane-changing feedback control for efficient lane assignment at motorway bottlenecks," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2625, no. 1, pp. 20–31, Jan. 2017.

[20] M. Ramezani and E. Ye, "Lane density optimisation of automated vehicles for highway congestion control," *Transportmetrica B, Transp. Dyn.*, vol. 7, no. 1, pp. 1096–1116, Dec. 2019.

[21] C. Roncoli, N. Bekiaris-Liberis, and M. Papageorgiou, "Optimal lane-changing control at motorway bottlenecks," in *Proc. IEEE 19th Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2016, pp. 1785–1791.

[22] T.-S. Dao, C. M. Clark, and J. P. Huissoon, "Optimized lane assignment using inter-vehicle communication," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2007, pp. 1217–1222.

[23] P. Varaiya, "Smart cars on smart roads: Problems of control," *IEEE Trans. Autom. Control*, vol. 38, no. 2, pp. 195–207, Feb. 1993.

[24] K. Kim, J. V. Medanić, and D.-I. Cho, "Lane assignment problem using a genetic algorithm in the automated highway systems," *Int. J. Automot. Technol.*, vol. 9, no. 3, pp. 353–364, Jun. 2008.

[25] S. Choi and H. Yeo, "Framework for simulation-based lane change control for autonomous vehicles," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2017, pp. 699–704.

[26] L. C. Edie, "Discussion of traffic stream measurements and definitions," in *Proc. 2nd Int. Symp. Theory Traffic Flow*, London, U.K., 1963, pp. 139–154.

[27] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 62, no. 2, pp. 1805–1824, Aug. 2000.

[28] J. L. Catbagan and H. Nakamura, "Desired speed distributions on two-lane highways under various conditions," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2088, no. 1, pp. 218–226, Jan. 2008.

Niharika Mahajan received the B.E. degree in civil engineering from the Birla Institute of Technology and Sciences, Pilani, India, in 2011, and the M.Sc. degree in transport and planning from the Delft University of Technology, The Netherlands, in 2015. She pursued research with the Delft University of Technology on designing algorithms for traffic control with intelligent vehicles to improve traffic performance. She is currently a Consultant in the domain of supply chain planning and optimization.

Andreas Hegyi received the M.Sc. degree in electrical engineering and the Ph.D. degree from TU Delft, The Netherlands, in 1998 and 2004, respectively. He is currently an Assistant Professor with TU Delft. He is the author or coauthor of over 100 papers. His research interests include traffic flow modeling and control, connected and cooperative vehicles, traffic state estimation, and traffic data analysis. He is a member of IEEE-ITSS and IFAC-CTS. He has served as the Program Chair for the IEEE-ITSC 2013 Conference, the General Chair for the IXth TRISTAN Symposium 2016, and an IPC member for various other conferences. He is an Associate Editor of IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS.

Serge P. Hoogendoorn received the M.Sc. degree in applied mathematics and the Ph.D. degree in civil engineering from the Delft University of Technology in 1995 and 1999, respectively. He is currently a Full Professor with the Delft University of Technology, where he holds the chair position in traffic operations and management and is a Distinguished Professor in smart urban mobility. He has authored or coauthored over 280 papers. His research interests cover a variety of topics, including traffic flow theory, traffic management, ITS, and active mode mobility. He is an IAC Member of the ISTTT. He is a member of the TRB Traffic Flow Theory Committee. He was a recipient of the ERC Advanced Grant in 2014. He is the Chair of the TRB Crowd Modeling and Management Subcommittee. He is an Editor of the *Journal of Advanced Transportation*, *EJTL*, and *EJTIR*.

Bart van Arem (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees in applied mathematics from the University of Twente, Enschede, The Netherlands, in 1986 and 1990, respectively. From 1992 to 2009, he was a Researcher and the Program Manager with TNO, working on intelligent transport systems, in which he has been active in various national and international projects. From 2009 to 2018, he was a Chair Professor in transport modeling with the Department of Transport and Planning, Delft University of Technology, The Netherlands, where he is currently a full-time Professor, focusing on the impact of intelligent transport systems on mobility. His research interests include transport modeling and intelligent vehicle systems.