Identifying Lane-Change Maneuvers with Probe Vehicle Data and an Observed Asymmetry in Driver Accommodation

Yiguang Xuan[1] and Benjamin Coifman, M.ASCE[2]

Abstract: This paper uses an instrumented probe vehicle to monitor ambient traffic and overcome many challenges of observing traffic flow phenomena that occur over extended distances. One contribution of this paper is a general methodology to identify the probe vehicle's lane of travel without a priori knowledge of where the lanes are. This knowledge is used to find the probe's lane-change maneuvers (LCMs), to differentiate these LCMs from GPS errors, and, in conjunction with a ranging sensor, to identify which lanes the ambient vehicles are in to find their LCMs. The second contribution of this paper comes from the identified LCMS. The data are used to provide an independent validation of earlier studies, and thus yield further evidence of how LCMs contribute to the formation of disturbances within freeway queues. In particular, it is found that vehicles following an entering vehicle generally complete their response and return to steady state quicker than those following an exiting vehicle. As discussed herein, this asymmetry in the lane-change maneuver accommodation time effectively induces a ripple in the traffic state that propagates upstream. The resulting disturbances provide a possible mechanism to explain the fact that congested traffic tends to fluctuate, e.g., stop-and-go traffic, rather than remain at a single, relatively stable congested state. **DOI: 10.1061/(ASCE)TE.1943-5436.0000401.**

CE Database subject headings: Highways and roads; Driver behavior; Vehicles; Traffic management; Traffic flow; Asymmetry.

Author keywords: Highway traffic; Lane-change maneuvers; Instrumented probe vehicles; Lidar; Traffic instabilities.

Introduction

In this paper an instrumented probe vehicle is used to monitor ambient traffic and overcome many of the challenges of observing traffic flow phenomena that occur over extended distances. As will be discussed shortly, one contribution of this paper is a general methodology to identify the probe vehicle's lane of travel without a priori knowledge of where the lanes are. This knowledge is used to find the probe's lane-change maneuvers (LCMs), to differentiate these LCMs from GPS errors, and, in conjunction with a ranging sensor, to identify which lanes the ambient vehicles are in to find their LCMs.

The second contribution of this paper comes from exploiting these new tools. The identified LCMs are used to provide an independent validation of Wang and Coifman (2008), and thus yield further evidence of how LCMs contribute to the formation of disturbances within freeway queues. This objective is important, because classical hydrodynamic traffic flow theory (Lighthill and Whitham 1955; Richards 1956) does not offer any mechanism as to how or why the commonly observed stop-and-go traffic forms. As asserted by Ahn and Cassidy (2007), the formation of the disturbances cannot be captured strictly via car-following in the absence of LCM. But the study of LCMs is complicated by the spatial nature of the maneuvers, the low density of conventional traffic detectors, and the large number of vehicles that travel in a freeway lane. These challenges are addressed by using the instrumented probe vehicle to monitor the ambient traffic to observe the traffic flow phenomena.

Among the body of research related to LCMs, some papers focus on the macroscopic properties, such as the fraction or frequency of LCMs (e.g., Kang and Chang 2004; Sheu 1999; Sheu and Ritchie 2001). Some papers approach the topic from the microscopic view, for example, by studying the distributions of time headways (e.g., Nakatsuji et al. 2006), or by applying gap acceptance models to microscopic traffic simulation (e.g., Gipps 1986; Yang and Koutsopoulos 1996; Ahmed et al. 1996; Zhang et al. 1998; Toledo et al. 2003; Hidas 2002, 2005). But there are relatively few papers on how LCMs impact the traffic state or actually cause delays to the traffic. Among these papers, Coifman et al. (2006) develop a model to estimate the delay caused by LCMs within a given lane relative to the situation in which no LCMs had taken place. Laval and Daganzo (2006) propose a model to explain the drop in discharge rate at bottlenecks due to LCMs and provide several simulations that appear to replicate empirical results observed from fixed-point detectors in earlier studies. Ahn and Cassidy (2007) examine the impacts of vehicles entering a given lane and how these vehicles contribute to the formation and growth of disturbances within a queue. They focused only on the entered lane, without considering the benefits to the exited lane or the combined impacts across the two lanes. Wang and Coifman (2008) used a set of complete vehicle trajectory data extracted from video over a short stretch of roadway to examine the mechanism underlying the delays in Coifman et al. (2006).

[1]Graduate Student Researcher, Institute of Transportation Studies, Univ. of California, Berkeley, CA 94720. E-mail: xuanyg@berkeley.edu

[2]Associate Professor, The Ohio State Univ., Joint Appointment with the Dept. of Civil, Environmental, and Geodetic Engineering, and the Dept. of Electrical and Computer Engineering, Hitchcock Hall 470, 2070 Neil Ave., Columbus, OH 43210 (corresponding author). E-mail: Coifman.1@osu.edu

Note. This manuscript was submitted on May 13, 2010; approved on December 29, 2011; published online on July 16, 2012. Discussion period open until January 1, 2013; separate discussions must be submitted for individual papers. This paper is part of the ***Journal of Transportation Engineering***, Vol. 138, No. 8, August 1, 2012. ©ASCE, ISSN 0733-947X/2012/8-1051–1061/$25.00.

J. Transp. Eng. 2012.138:1051-1061.

Wang and Coifman show that the impacts of lane-change maneuvers are not balanced: vehicles following an entering vehicle generally complete their response and return to steady state quicker than those following an exiting vehicle. This lane-change maneuver accommodation time (LCAT) imbalance propagates upstream, and it appears to be a source of speed and flow fluctuations (or oscillations) within a queue. Although both lanes behind an LCM undergo a disturbance that propagates upstream, the two lanes do so out of phase with one another. While the net gain in flow in the exited lane balances the net loss in flow in the entered lane over time, the latter does so over a shorter period of time and thus undergoes a larger displacement in instantaneous flow. So at any given point in time and space the impacts in the two lanes do not cancel each other; thereby creating a ripple that propagates upstream. But Wang and Coifman used trajectory data from a single lane at a single facility (in the outside lane immediately upstream of an on ramp), and they used only one hour of data collected on a single day (all that was available at the site). The present work seeks to demonstrate the same LCAT imbalance over a more diverse data set. Ultimately the empirical results from these studies should help develop more robust microscopic LCM models that better capture the impacts of LCMs on traffic.

As noted above, an instrumented probe vehicle is employed to monitor the ambient traffic. The present work seeks to identify the LCMs and classify them using the probe vehicle trajectories [from differential global positioning system (DGPS)] and the ambient traffic [via light detection and ranging (LIDAR) scans]. Key to this effort is establishing the appropriate references: (1) a reference trajectory in space to establish where the lanes are relative to the vehicle's current location, an important step to identify when LCMs occur, and (2) the set of quasi-equilibrium states forming the speed–spacing relationship, this set of states is then used to identify the start and end of LCAT. The details of these processes are presented in the following sections.

There are many examples where additional sensors are used on a vehicle to detect threats to that vehicle either for collision avoidance, driver behavioral studies, or for autonomous vehicle control, but these systems generally are not employed for traffic flow theory development. Traditionally traffic flow theory has been developed using fixed-point detectors, with no direct information about conditions between the detector stations. More recently a small number of vehicle trajectory data sets have been collected over short stretches of roadway (on the order of 0.5 km) for a short duration (on the order of an hour); most notably the next generation simulation (NGSIM) effort, (FHWA 2006a, b) and the Turner Fairbanks data sets (Smith 1985; Smith and Mark 1985). The instrumented probe vehicle falls somewhere in between: providing both the extensive spatial coverage of a network of fixed-point detectors and the rich details of vehicle trajectories between the fixed points, but only for the ambient traffic around the probe vehicle's own trajectory. Since the probe vehicle passages are separated by hours, days, or weeks, compared to NGSIM, it is much less likely that all of the observations could be influenced by a single confounding factor (e.g., weather or an incident). The three different approaches are in fact a strong complement to one another, each with its own strengths.

Overview

The remainder of this paper is organized as follows: the first section provides a data description, including the probe vehicle sensors, explanations of the data available, the routes, and other details. The second section presents the tracking process for the LIDAR data. The LIDAR data are rudimentary, merely providing the distance to the nearest object at half-angle increments. So this section provides an overview of how individual vehicles were segmented from the background in the LIDAR data and then trajectories for these vehicles were extracted relative to the probe vehicle coordinate system. Unfortunately, if the probe vehicle undertakes an LCM, the sensors' frame of reference moves, and from the sensor data it will look as if all of the tracked vehicles moved to the opposite direction. So the third section describes how a reference trajectory was generated and used to identify LCMs both by the probe vehicle and the ambient traffic, without a priori knowledge of the lane locations. The fourth section then defines the process underlying the lane-change accommodation time calculation via a set of reference quasi-equilibrium states representing drivers' preferred speed–spacing relationship. The LCAT is then calculated for many maneuvers in congested conditions (with speed less than 72 km/h during the LCM). The results show asymmetry in the LCAT when the probe vehicle follows LCMs by entering and departing lead vehicles. The asymmetry in LCAT also holds when the probe vehicle undertakes LCMs, either decreasing or increasing the relative spacing as it changes lanes and thus, changing lead vehicles. Finally, the paper closes with the conclusions of this work.

Data Description

Sensors

A van equipped with multiple sensors is used as the probe vehicle for data collection, shown in Fig. 1. As noted on the figure, there are five types of sensors installed on the van. Only two of the sensors are used in this research, namely: the forward-facing LIDAR, and the DGPS. For validation purposes, there is also a camera to capture 320×240 pixel digital images of the forward view at 1 Hz.

The forward-facing LIDAR uses a laser beam to measure the distance to surrounding objects. It scans continuously from in a horizontal plane, with a frequency of 3.3 Hz, an angular coverage of 180°, and a 0.5° angular resolution. The range of the LIDAR sensor is 81.91 m, with a resolution of 0.01 m.

The DGPS receiver used in this work is a Trimble AG132 GPS receiver with Omnistar VBS corrections. It is an L1 only (single frequency) receiver with 12 channels. Omnistar VBS corrections are processed in real time. According to the receiver specifications, the DGPS data are accurate to within 1 m for 95% of the time. The 1-Hz DGPS data includes the following information: time stamp (seconds after midnight), latitude (degrees), longitude (degrees), velocity (meters/second), heading (radians), differential status, and altitude (meters). The DGPS data might include large transient errors (compared to the resolution) due to occlusion or multipath, which usually happens as the vehicle passes under an overpass or other occluding roadside feature, as discussed below in the section about identifying lane-change maneuvers.

Routes and Other Details

The driver of the probe vehicle is instructed to drive on one of two different routes in Columbus, Ohio. Both routes begin in the central business district (CBD) and head north along I-71. The first route is termed the travel-time route, in which the driver completes two 50-km round trips from SR-315 to Polaris Pkwy on I-71 (each round trip termed a travel-time run). The driver is instructed to drive in the second lane from the center of roadway, except when choosing to overtake the vehicle ahead of him or her. After preliminary analysis, it was found that the section of I-71 from North Broadway to Polaris Parkway is typically free flowing, so a new route was deployed in the course of this research to focus on

J. Transp. Eng. 2012.138:1051-1061.

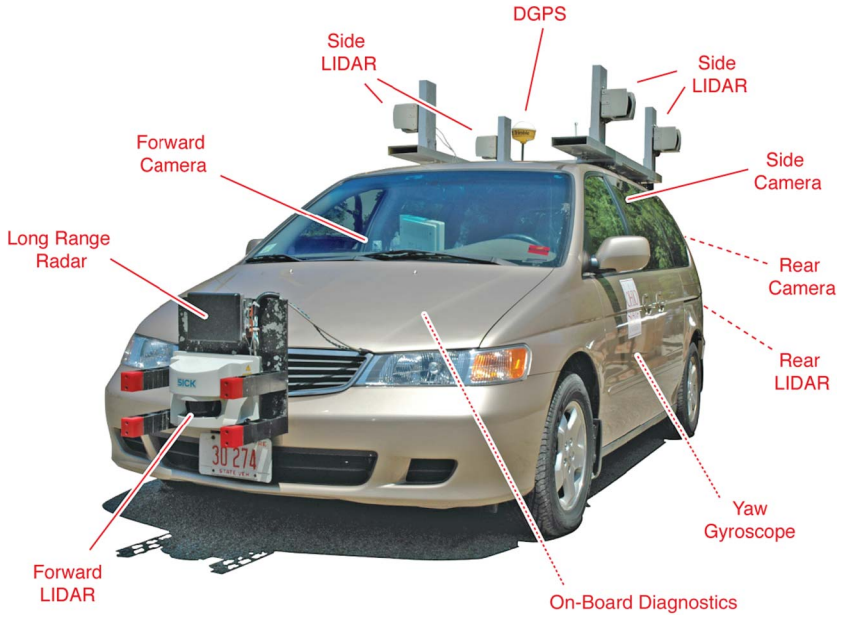

Fig. 1. Instrumented probe vehicle with the various sensors highlighted

the segments where congestion is most common. Termed the free-style route, this new route differs from travel-time route in two important ways. First, the driver will complete three 23-km round trips from SR-315 to North Broadway on I-71 (similarly, each round trip termed a free-style run). Second, because the LCM behavior on the road is of particular interest, the driver is free to choose any lane at any time, hence the term "free-style".

All of the data collected on a given tour of a route comprise a single data set. Data were collected between June 2005 and August 2006. A total of 29 travel-time route data sets and 16 free-style route data sets are used for this research. So the portion of I-71 between the CBD and North Broadway is observed a total of 106 times in each direction because there are two round trip runs in each travel-time route, and three round trip runs in each free-style route. The data were collected by six different undergraduate student drivers, but no distinction is made among the drivers in the analysis.

Tracking

The LIDAR provides rich information about the surrounding vehicles. Every 0.3 s, the LIDAR scans the surrounding objects in a plane roughly 0.5 m above the ground. The range and angular information produces a 2D image of the position of the nearest object (within the range of the sensor) at each angle sampled. Such an image is called a frame throughout the rest of the paper. To illustrate this process, Fig. 2(a) shows a hypothetical top-down view of the roadway with the instrumented vehicle shaded at the bottom. The LIDAR sweeps 180°, at 0.5° increments [Fig. 2(b)], and receives returns from vehicles and fixed objects [Fig. 2(c)]. Finally, Fig. 2(d) shows the resulting frame of data returned from the LIDAR scan relative to the vehicle coordinates.

Ultimately the goal is to track the distinct vehicles throughout the duration that they reside in the LIDAR field of view, thereby producing vehicle trajectories relative to the probe vehicle's trajectory. A given target will appear differently in the frame depending on the relative position of the object to probe vehicle, and it is not always immediately apparent whether a target is a vehicle or a stationary object, both may take on a similar appearance. The vehicles need to be segmented from stationary objects and from one another. Based largely on the work by Wang and Coifman (2005, unpublished internal report), Gao and Coifman (2006), and Gao and Coifman (2007), this tracking task is split into three components: the grouper, the classifier, and the tracker. Each component is described below. While the details of the basic tracking process can be found in the papers by Gao and Coifman, this section briefly reviews the tracking process. First, the grouper clusters the LIDAR data points of each frame into discrete objects based on the Euclidian distance between the data points. Thresholds in distance are set empirically to ensure that data points from the same object (vehicle or stationary objects) are usually grouped together, while also being segmented from all other discrete objects.

Next, the classifier examines each discrete object reported by the grouper in the frame. Using the shape and the history from preceding tours, a given object is classified to differentiate vehicles from roadside boundaries. The shape of a vehicle cluster can be a horizontal line segment, a vertical line segment, or a combination of the two in an *L* shape (or reverse *L* shape), depending on the relative position of the object to probe vehicle. Many stationary nonvehicle objects can take on these same appearances as the vehicles. The nonvehicle objects that are not readily distinct from vehicles can be distinguished if the history from many runs is employed, and the observed objects from the probe vehicle's coordinates are projected to the world coordinates. A given stationary nonvehicle object will be observed at the same location on all runs in which it is not occluded by vehicles, while the vehicles will be observed at random locations on the roadway with a much lower density. After many runs, the locations with stationary objects will have a high frequency of observations, while the intervening locations with vehicles will have a lower frequency of observations. Gao and Coifman (2007) use this fact to identify the regions that are on the road. If a cluster is on the road it is considered to come from a vehicle and ignored otherwise.

The on-road clusters are considered to be vehicles and only these clusters are tracked. A Kalman filter is used to model the 2D vehicle position relative to the LIDAR sensor [i.e., the two axes in Fig. 2(d)], assuming constant relative speed and it is then used to estimate the vehicle position in the next frame. The association of vehicles between frames is based on the Euclidian distance between the estimated position from the Kalman filter and the measured

J. Transp. Eng. 2012.138:1051-1061.

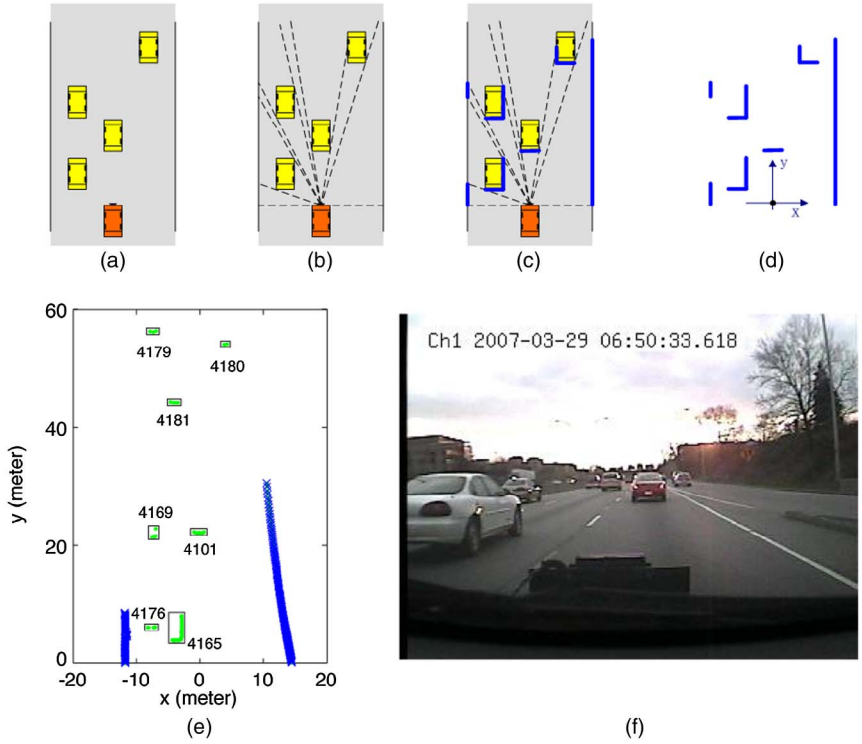

Fig. 2. Example of LIDAR detection (a) top-down view of the roadway, the instrumented vehicle is shaded at the bottom; (b) forward LIDAR scans the world; (c) receiving returns from vehicles and fixed objects; (d) the resulting frame of data returned from the LIDAR scan relative to the vehicle coordinates; (e) an actual sample of the forward LIDAR "view" after segmenting vehicles in one frame; (f) the concurrent digital image

position of clusters observed in the current frame. The thresholds for distance are also set empirically. The methodology is able to extract information about the 3D trajectories (x, y, t) of the surrounding vehicles relative to the probe vehicle.

Here, Fig. 2(e) shows an example of the forward LIDAR "view" from the van and Fig. 2(f) the concurrent digital image from the camera. The LIDAR sensor is at $(0, 0)$. The targets have been grouped, as shown with boxes around the cluster of points from each vehicle; classified, as shown with points for vehicles and "x" for stationary background objects; and vehicles tracked, as indicated by a unique target number for each vehicle cluster. Henceforth an individual cluster will be called a target vehicle. The field of view of the camera is narrower than the LIDAR scan, so target vehicle 4,165 is evident on the left-hand side of the image.

Identifying Lane-Change Maneuvers

Moving beyond earlier tracking efforts, the present work seeks to explicitly detect lane-change maneuvers and measure the associated disturbances.

Identifying All the LCMs Relative to the Probe Vehicle

With the 3D trajectory information relative to the probe vehicle coordinates, the process of identifying LCMs among the target vehicles is conceptually simple. Namely, one can find when a target vehicle has a lateral displacement approximately equal to one lane width (3.6 m). While such a lateral displacement is indicative of an LCM, it will arise both when a target vehicle changes lanes and when the probe vehicle itself changes lanes. Simply put, an LCM by the probe vehicle in one direction will result in the apparent phantom LCM of all surrounding vehicles to the opposite direction in the LIDAR data. The probe vehicle's LCMs will be accounted for via the reference trajectory presented in the next section. At this first stage the simple displacement methodology has the following limitations.

1. When an LCM occurs, the relative motion information does not indicate which of the two sources occurred: the probe vehicle changed lanes or the target vehicle changed lanes.
2. When a tracked vehicle makes an LCM at the same time as the probe vehicle, it may appear as if no LCM occurred, but there are actually two LCMs (one by the probe vehicle and one by the target vehicle).
3. When the road merges (or diverges), vehicles coming from (going to) a different origin (destination) than the probe vehicle will exhibit non-LCM lateral motion. If care is not taken, their lateral motion may erroneously be attributed to LCMs when in fact their lane of travel is not parallel to the probe's lane of travel.

The first two points can be addressed by using the positioning data from the DGPS and other sensors to independently determine when the probe vehicle changes lanes, as will be done in the next section. The final point often occurs at ramps. The problem cannot be solved unless additional information about the roadway geometry is employed. Because the probe vehicle used in this study rarely travels in the outside lane, the impacts of the third point are mitigated by excluding any LCM that does not directly involve the probe vehicle. Thus, the scope is limited to LCMs that occur between the probe vehicle's current lane and an immediately adjacent lane. This step has the added benefit of ensuring that there is an unoccluded view of the lead vehicle during the LCM.

Establishing a Reference Trajectory and Identifying the Probe Vehicle LCMs

One challenge of this work is to identify LCMs without a priori knowledge of where the lanes are. To accommodate for the fixed geometry of the roadway multiple probe vehicle trajectories through the roadway segment are integrated to establish a reference

J. Transp. Eng. 2012.138:1051-1061.

trajectory in a single lane that is then used to identify the other discrete lanes, as discussed in the first subsection. To simplify the process this derivation uses only the travel time runs, because the drivers are proscribed to maintain a specific lane most of the time (this reliance on a dominant lane of travel can be dropped with only minor modification to the analysis). But the process is complicated by the fact that there are a few mandatory LCMs (MLCs) along the probe vehicle's route and thus, at these points the reference trajectory jumps from one lane to another if care is not taken to identify MLCs by the probe vehicle, as discussed in the second subsection. After controlling for the MLCs, the reference trajectory falls in a single lane throughout the run and travel in the other lanes is evident by a fixed lateral displacement by an integer number of lane widths. The revised reference trajectory is then used to identify discretionary LCMs (DLCs) by the probe vehicle as discussed in the third subsection. The reference trajectory is also used in subsequent sections to find LCM by the ambient vehicles.

Establishing the Reference Trajectory

The objective of this section is to establish a robust reference trajectory that defines a curvilinear coordinate system, with the abscissa corresponding to the lateral distance (across the road), and the ordinate corresponding to the longitudinal distance (along the road). This reference trajectory is built from many noisy individual trajectories recorded in the probe vehicle DGPS. The data used here are the travel-time run data sets, because the driver is instructed to stay in the second lane from the center of roadway except when overtaking. Thus, the trajectories should usually overlap in the same lane and most of the time a given travel-time run trajectory should fall within close vicinity of the reference trajectory, with occasional deviations arising from LCMs or GPS errors.

First, an arbitrary trajectory, say T_1, is chosen, and the points on T_1 are initially taken to be $(0, Y)$, where Y denotes the longitudinal distances along T_1. Next, the coordinates $[X_i'(Y), Y]$ of all the other trajectories T_i ($i = 2$ to n, where n is the total number of trajectories) are calculated by projecting them laterally onto T_1, where $X_i'(Y)$ denotes the lateral distance of T_i to T_1 at location Y. Next, the reference trajectory is defined as the median of the lateral distances of all trajectories at the given Y, $X''(Y) = \text{median}\, X_i'(Y)$. The reference trajectory is set to be $(0, Y)$ and the lateral distance to the reference trajectory, $X_i''(Y)$, is calculated for each run (so at this point, in general, T_1 will have a nonzero abscissa at a given Y). The median was used rather than the mean because the median is less sensitive to outliers in the dataset, e.g., the median will not be affected by occasional DLC while the mean would yield a reference trajectory that includes the impacts of every DLC. In the event that the probe vehicle was free to choose lanes, the methodology could be modified to use the modes of the lateral distribution instead of the median.

Identifying Mandatory LCMs by the Probe Vehicle

To measure lateral position across the roadway, it is necessary to correct for the mandatory LCMs. The MLCs occur when the probe vehicle has to shift lanes to follow the given route, e.g., because of geometric features. While the driver may need to change lanes for an MLC in every run, the exact location will vary from one run to another; i.e., the MLC will occur over a range of Y coordinates. Consider an MLC observed across many trajectories. One of the trajectories will begin the MLC farther upstream than all of the others. Moving downstream, more and more of the trajectories will change lanes until the last trajectory does so. As one progresses downstream through this window, more and more trajectories will shift away from the median used for the reference trajectory (in the direction of the MLC) until the reference trajectory jumps over to the new lane and the remaining trajectories now become prominent on the opposite side of the reference trajectory until reaching the end of the window. Compared to the reference trajectory, most individual probe vehicle trajectories will typically appear to make two LCMs. One of these LCMs is that individual trajectory's true MLC and the other is false LCM that actually captures the lateral jump in the reference trajectory. To illustrate this point, in the travel-time run data sets, around longitudinal distance 5 km, there is an MLC to the right (in this case caused by the combination of a lane drop and the driver's instruction to stay in the second lane). Here, Fig. 3(a) shows the lateral distance of all the travel-time run trajectories with respect to the reference trajectory, many of the trajectories make the MLC prior to 5.7 km as evidenced by an LCM to the right followed by the false LCM to the left when the reference trajectory changes lanes at 5.7 km. Other trajectories make the MLC after 5.7 km, so in these cases the false LCM to the left comes before the true MLC to the right. Finally a few of the trajectories make the MLC close to 5.7 km and show little evidence of any

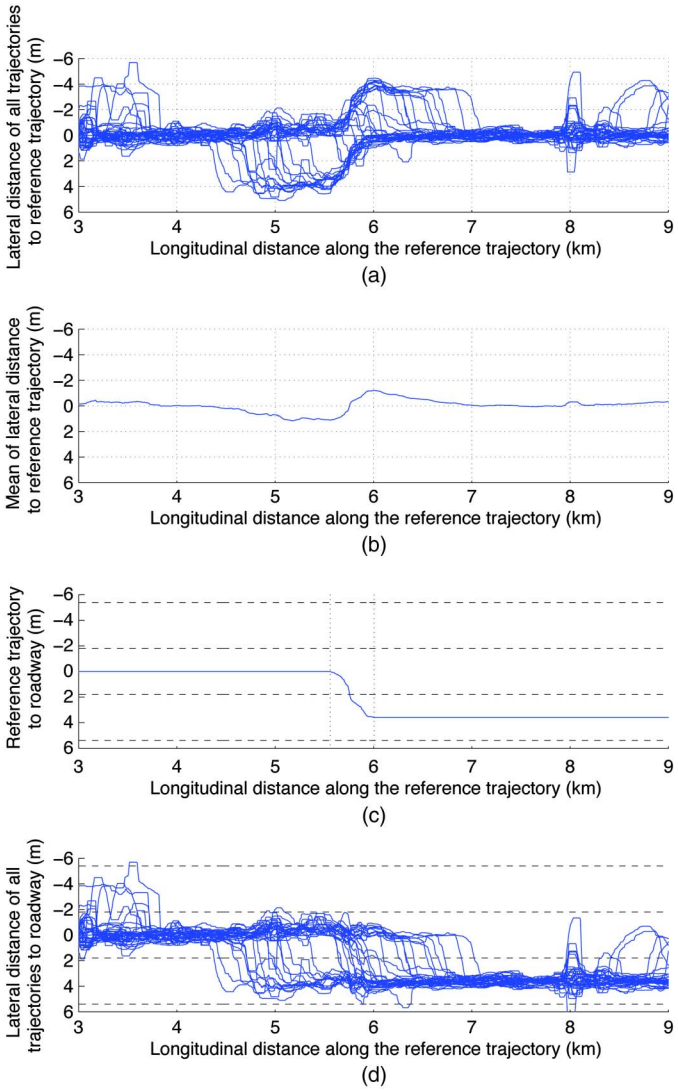

Fig. 3. Probe vehicle MLCs around longitudinal distance 4.3 km along the travel-time run, all runs travel in the direction of increasing longitudinal distance (a) lateral distance of all trajectories with respect to reference trajectory; (b) mean of lateral distances; (c) difference trajectory with respect to roadway; (d) all trajectories with respect to roadway after combining (a) and (c)

J. Transp. Eng. 2012.138:1051-1061.

LCM because the true and false maneuvers cancel one another in this plot.

Within the window, the individual MLCs disrupt the reference trajectory because it uses the median lateral position. Fortunately, the MLC will have a different impact on the mean of lateral position; rather than changing abruptly, the mean lateral position will gradually shift along the length of the longitudinal range in which the MLC fall. The difference between the mean and median is used to identify locations of MLCs in the reference trajectory. This difference will shift first in the direction of an MLC and then when the median shifts lanes, the difference will jump to the opposite side of the reference trajectory. Here, Fig. 3(b) shows the mean of lateral distances relative to the reference trajectory, and as expected, the mean first drifts to the right and then crosses zero and jumps to the left when the reference trajectory changes lanes to the right at roughly 5.7 km.

Although the reference trajectory jumps lanes over a short longitudinal distance, it does not do so instantaneously and so this work sought to capture its progression. Almost all of the trajectories exhibit the same false LCM to the left due to the reference trajectory actually changing lanes to the right, the only exceptions being those few trajectories that make the LCM concurrent with the reference trajectory. But in this reference plane, the individual trajectories that made the maneuver prior to 5.7 km are one lane width (3.6 m) below the trajectories that do so after [the two dense regions in Fig. 3(a) at 5.7 km]. So 3.6 m is added to those trajectories that make an LCM prior to the zero crossing in Fig. 3(b). The mean lateral position across all of the trajectories (including those shifted 3.6 m) is then subtracted from the reference trajectory, resulting in Fig. 3(c). Throughout the remainder of this paper, the reference trajectory is assumed to incorporate this MLC correction unless explicitly noted otherwise. In Fig. 3(d) the corrected trajectories are shown with respect to the roadway. Each trajectory now exhibits a single MLC (to the lane on the right) without any of the phantom LCMs due to the reference trajectory changing lanes. After subtracting out the shift in the reference trajectory, the exact locations of the MLCs in a given data set can be found using the same techniques, as presented in the next subsection to identify discretionary LCMs.

Identifying Discretionary LCMs by the Probe Vehicle

Given the reference trajectory and a specific probe vehicle trajectory, the lateral distance to the reference trajectory is found, e.g., as shown in Fig. 3(d). During a discretionary LCM (DLC), the probe vehicle should be offset laterally by a lane width, which is roughly 3.6 m. So threshold lines are set with lateral distance 1.8 m from the reference trajectory. The threshold lines correspond roughly to the lane lines, and whenever a trajectory crosses any threshold line it is considered a DLC.

Most of these DLCs in the travel-time run data sets are from the driver overtaking another vehicle and then returning to the original lane. An overtaking will usually show up as a lateral deviation beyond a threshold line, and then return back to the original lane after some time, e.g., as shown in Fig. 4. But not all of the lateral deviations beyond a threshold are due to DLCs; some disturbances come from GPS errors due to obstructions and multipath [e.g., one can see disturbances around 8 km in Figs. 3(a) and 3(d) that arise from an overpass]. Fortunately, most of these GPS positioning errors are large in magnitude but short in duration, e.g., while reacquiring a lock on the satellites during one or two samples after emerging from an underpass. Such short transient errors can be quickly filtered out using a moving median (e.g., as per Coifman and Dhoorjaty 2004) on the time-series lateral distance from the reference trajectory. In contrast, a real overtaking maneuver will usually take longer. So the out-of-threshold-line time is calculated whenever a trajectory is beyond the first lateral threshold line.

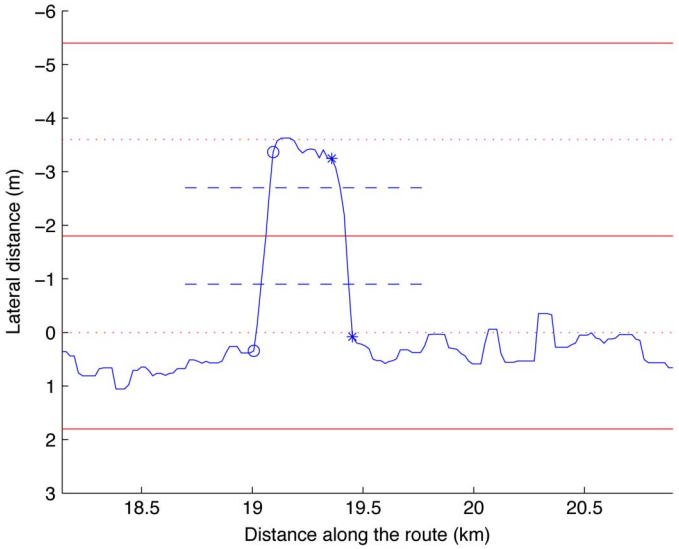

Fig. 4. An example of the probe vehicle executing an overtaking maneuver, comprised of two successive LCMs

The camera imagery was used to verify the source of all departures from the lane, so as to differentiate between an overtaking and a disturbance. Here, Fig. 5 shows the cumulative distribution function (CDF) of the out-of-threshold-line time. Based on the manual verification of 30 actual overtaking maneuvers and 57 disturbances, most of the overtaking maneuvers can be differentiated from the disturbances simply from a minimum out-of-threshold-line time. No overtaking is missed if the time threshold is set to 10 s. Assuming these data are representative, two successive DLCs in opposite directions will not typically occur within 10 s.

Across the data set there are two GPS errors that are erroneously accepted as DLC by this simple filter. The filter assumes a straight line trajectory between successive GPS points, without accounting for the time step between GPS observations. However, on two passes the GPS dropped out while the van was on a curve and

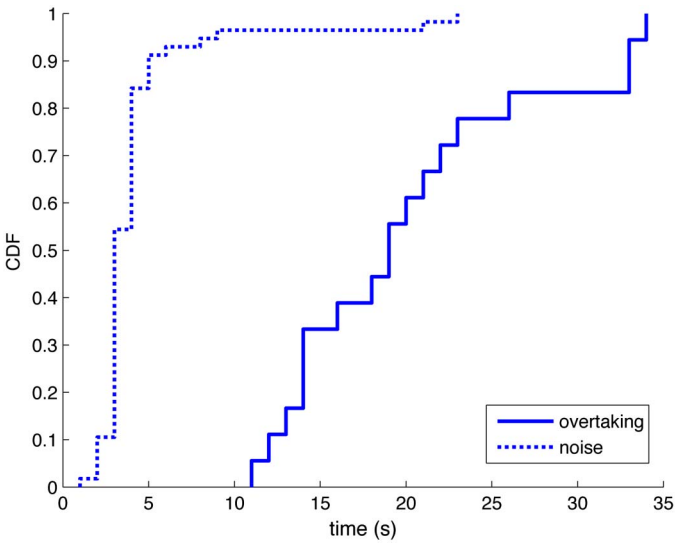

Fig. 5. CDF of the out-of-threshold-line time for overtaking and GPS disturbance

J. Transp. Eng. 2012.138:1051-1061.

the straight line trajectory assumption resulted in a large lateral deviation from the reference trajectory. For the present study, the two DLC errors are excluded from further consideration and the manual classifications from Fig. 5 are used. In general, this problem can be addressed by suppressing any possible DLC that occur while the GPS is momentarily unavailable.

Returning to the Target Vehicle Trajectories

After establishing a probe vehicle reference trajectory with corrections for MLCs and accounting for any DLCs in the specific probe vehicle trajectory, there is now (X, Y, t) from the probe vehicle run. The 3D trajectories of the surrounding vehicles measured relative to the probe vehicle (x, y, t) are then projected to their physical location along the road, by taking the sums, $X(t) + x(t)$, and $Y(t) + y(t)$. Fig. 6 shows two examples of target vehicle LCMs. In both cases the probe vehicle did not change lanes and although not shown, the probe vehicle is located at $[X(t), 0]$ on the plots. In each case the trajectory shows $[X(t) + x(t), Y(t) + y(t)]$ for the target vehicle. In Fig. 6(a), a vehicle in the left lane changes into the current lane of the probe vehicle. In Fig. 6(b), a vehicle in the right lane first enters the probe vehicle's current lane and then continues to the left lane.

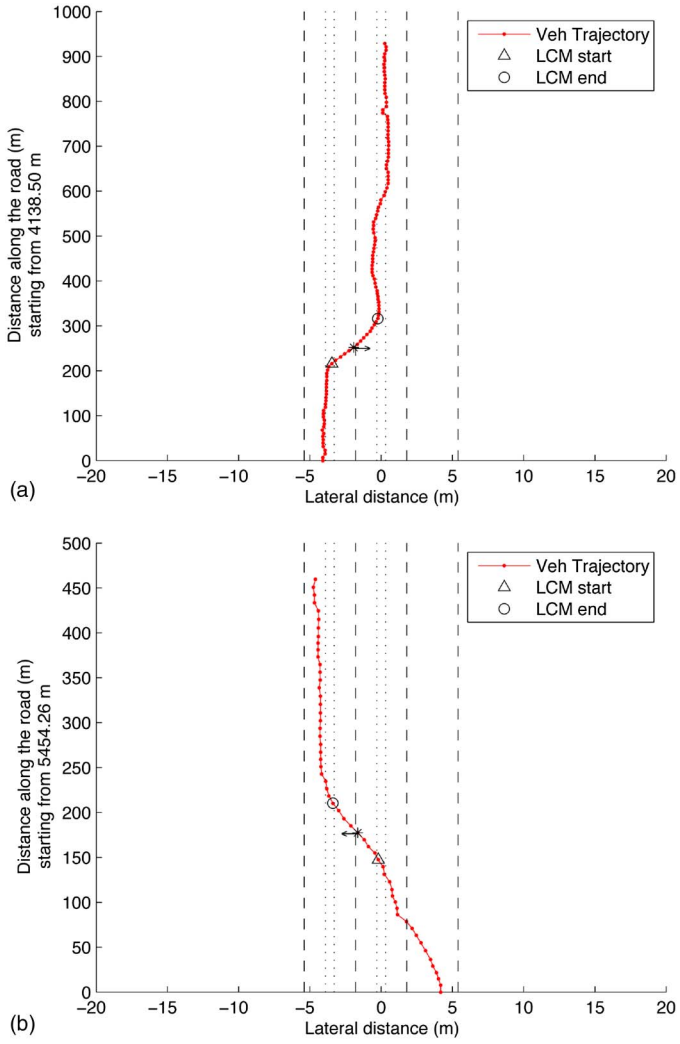

Fig. 6. Two examples of finding LCM of an ambient vehicle based on LIDAR tracking results: (a) from the left lane; (b) to the left lane (the earlier LCM from the right lane is not highlighted); lateral distance is relative to the probe vehicle's lane of travel; the arrows indicate the direction of the LCM, long dashed lines show the lane lines, and short dashed lines show the center of lane with a tolerance of 0.3 m

Lane-Change Accommodation Time

Much of traffic flow theory is built upon the fundamental relationship between speed, flow, and density; or alternatively in the context of car following, between speed, headway, and spacing. In either case, the traffic state (i.e., the three parameters) is typically assumed to fall on or near a single curve, e.g., flow versus density or speed versus spacing relationships. Shock waves and other disturbances can cause a transient deviation away from the curve, but the traffic state quickly returns to the curve. Much of the existing body of traffic flow theory ignores LCMs, assuming the impact is negligible. But when an LCM occurs, the spacing of several vehicles will abruptly change, and each involved vehicle will have to adjust their speed to return to the driver's preferred speed–spacing relationship. This lane-change accommodation time does not occur instantaneously, so each LCM will perturb the traffic state for a short time, and thus, also perturb the vehicle's trajectory. Since empirically observed speed–spacing relationships are noisy, as will be discussed shortly, one has to use thresholds of some form to define when a driver is within his or her steady state speed–spacing for the given conditions.

Under ideal conditions an LCM in congestion will not reduce the net flow or increase the net delay across the two lanes, but one lane benefits at the expense of the other for a short duration, proportional to the LCAT. As shown in Wang and Coifman (2008), the fact that the LCM disrupts the lead vehicle trajectory in a given lane means the following vehicles must also follow the perturbation, disrupting the traffic state in that lane and the net result will be manifest as a ripple propagating upstream. This provides one source for disturbances to form in queues and potentially be a source of unstable stop-and-go traffic. It is also important to note that the LCAT is experienced by the drivers as they travel down the road. The resulting disturbances propagate upstream and when they are viewed from a stationary location on the side of the road, the duration of the impact of a given disturbance is longer than that experienced by the driver (being a function of the vehicle speed and the speed that signals propagate upstream). Wang and Coifman found that the LCAT was imbalanced between the exited lane and the entered lane; thus, the LCM also induces a ripple in the traffic state when summed across lanes. But Wang and Coifman only used data from one hour, in one lane, on one facility. This section seeks to provide an independent validation of the LCAT imbalance at other locations, using the instrumented probe vehicle data.

This work examines the speed–spacing relationship from the probe vehicle, where spacing is defined from the rear bumper of the lead vehicle to the rear bumper of the probe vehicle. Of course one must define a preferred speed–spacing relation before being able to detect deviations from it. In reality the speed–spacing data are very scattered, so the first subsection defines the quasi-equilibrium state to determine when a driver begins and ends their accommodation to an LCM. The second subsection develops the lane-change accommodation process, starting when the time-series speed–spacing relation departs from the defined quasi-equilibrium state before an LCM, and ending when the time-series speed–spacing relation first returns to the quasi-equilibrium state after the LCM. The third subsection presents the results and analysis of the measured LCATs.

J. Transp. Eng. 2012.138:1051-1061.

Quasi-Equilibrium State

This section defines quasi-equilibrium state in the speed-spacing plane. To address the fact that the speed–spacing relationship may vary from data set to data set, all of the individual speed–spacing measurements from all of the available data sets are plotted together, yielding the large cloud of points in Fig. 7(a). In Fig. 7(b), the corresponding density of the data points is shown, where unit density is defined as the density that would be observed if the data were uniformly distributed over the region of speed–spacing plane shown in the plot.

Next, segmenting the data in to speed bins every 3.6 km/h, the spacing distribution is evaluated in each bin. The following percentiles of spacing are calculated for each speed bin: 30%, 35%, 40%, 60%, 65%, and 70%, and are shown with the curves in Fig. 7(c). The first Highway Capacity Manual (BPR 1950) employed the results from 23 studies conducted between 1924 and 1941 that examined the speed–spacing relation for the purpose of estimating capacity. Among the 23 studies, 22 adopted a speed–spacing relationship in the form of a second-order polynomial

$$S = \alpha + \beta V + \gamma V^2 \quad (1)$$

were S = spacing and V = speed. The parameters have specific interpretations: α = the effective vehicle length: β = the reaction time; and γ = the reciprocal of twice the maximum average deceleration of the following vehicle. Although the origins are more than 50 years old, this model is still often used today (e.g., Rothery 2001). Borrowing this framework, the percentile curves from Fig. 7(c) are smoothed via a second-order polynomial linear regression, as shown in Fig. 7(d). The R^2 value is at least 0.98 for each of the six fitted polynomial curves.

The smoothed 35th and 65th percentile curves are taken as the bounds of the quasi-equilibrium state at the given speed. The choice of 35th and 65th percentile curves to define the quasi-equilibrium state was somewhat arbitrary. These bounding percentile curves were chosen so that for most of the time in car following (i.e., away from any LCM) the speed–spacing relation will lie within the quasi-equilibrium bounds, and when an LCM occurs, the deviation in spacing will be large enough to exceed the quasi-equilibrium bounds. The other percentiles are used for sensitivity analysis to

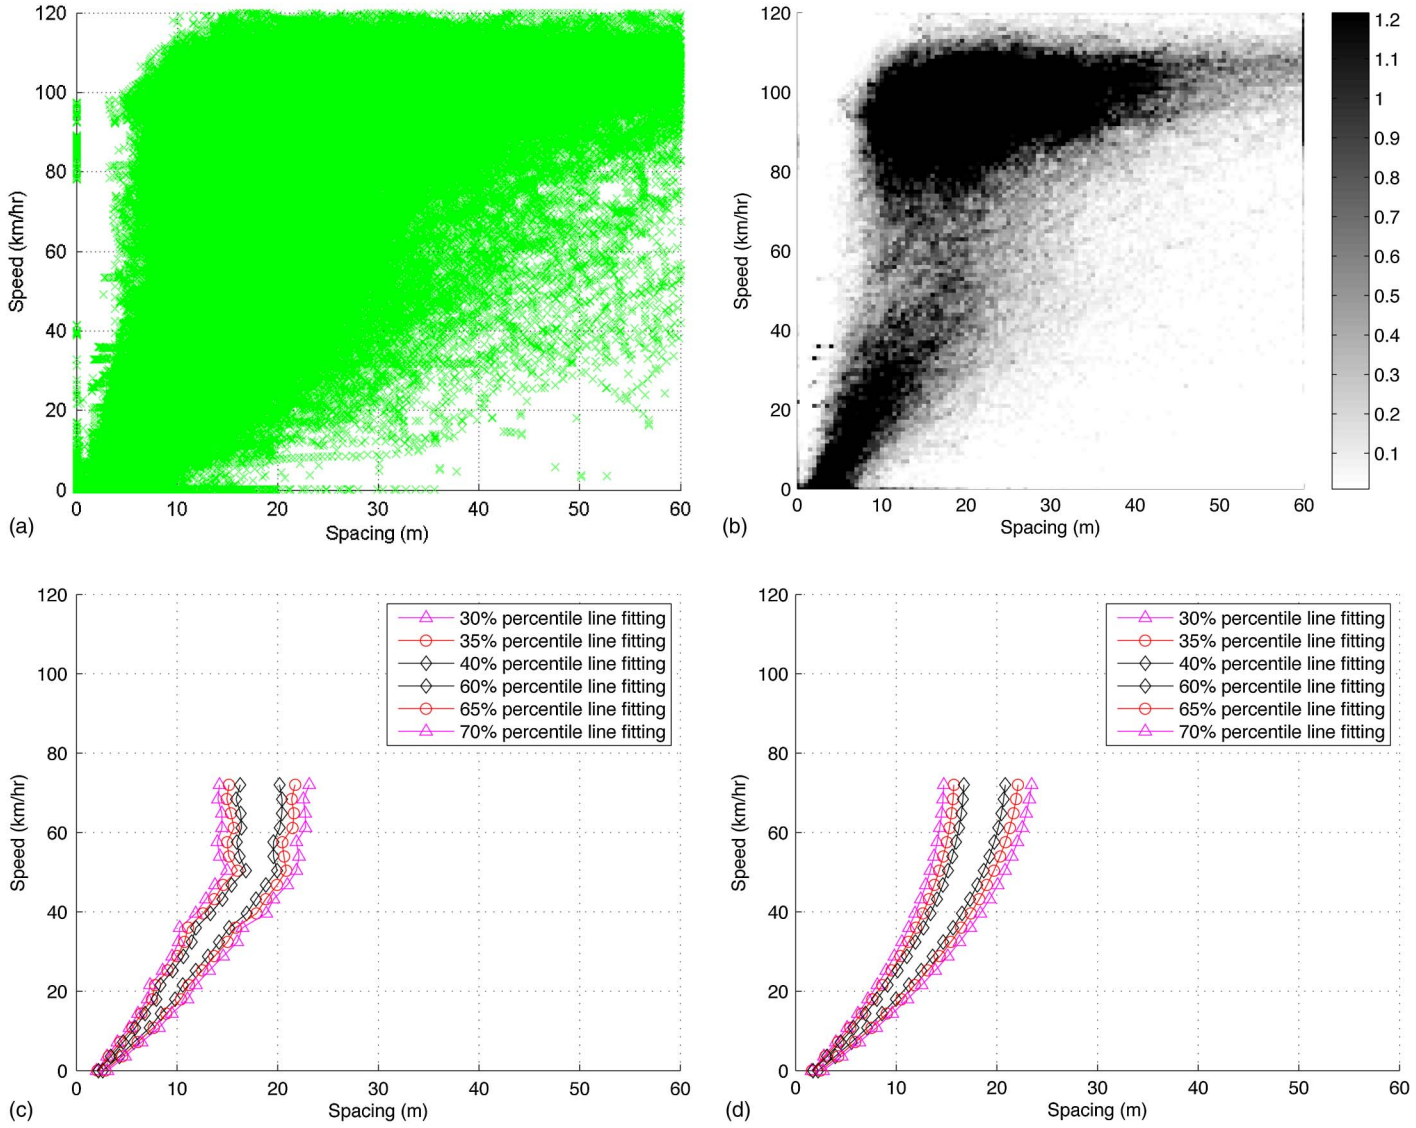

Fig. 7. Speed–spacing relationship from the probe vehicle: (a) all data on all runs; (b) density plot showing the density of the data normalized with respect to the average density; the resulting percentile curves across each speed bin; (c) raw; (d) after second-order polynomial fitting

J. Transp. Eng. 2012.138:1051-1061.

ensure there are no significant impacts arising from the choice of specific percentiles.

Of course this approach involves other trade-offs. Presumably the data in a given run would yield a tighter range, e.g., the data in Fig. 7 come from six different drivers. But this approach to calculate the percentile lines requires many samples in each bin. So the specificity of a given run was sacrificed for the benefit of the much larger sample size of the entire data set; this fact is particularly important for bins that have few observations in a given run.

Lane-Change Accommodation

With the quasi-equilibrium state one can now identify periods when the driver deviates from the preferred speed–spacing relation. The quasi-equilibrium state is employed solely to decide when a driver begins and ends his or her accommodation to an LCM. In turn, the lane-change accommodation is used to compare the behavior of drivers behind a vehicle that enters and a vehicle that departs the lane. As illustrated below, the specific values of the LCAT in this study are highly dependent on many parameters, and they are only meant for relative comparisons when those parameters are held constant. The LCAT is defined as beginning when the speed–spacing relation last leaves the quasi-equilibrium state (crossing a bounding percentile line) immediately prior to the LCM and lasts until the speed–spacing relation first returns to the quasi-equilibrium state (again crossing a bounding percentile line) immediately after the LCM. During the accommodation process, the driver following the maneuver adjusts speed in an effort to return to quasi-equilibrium state behind the new lead vehicle.

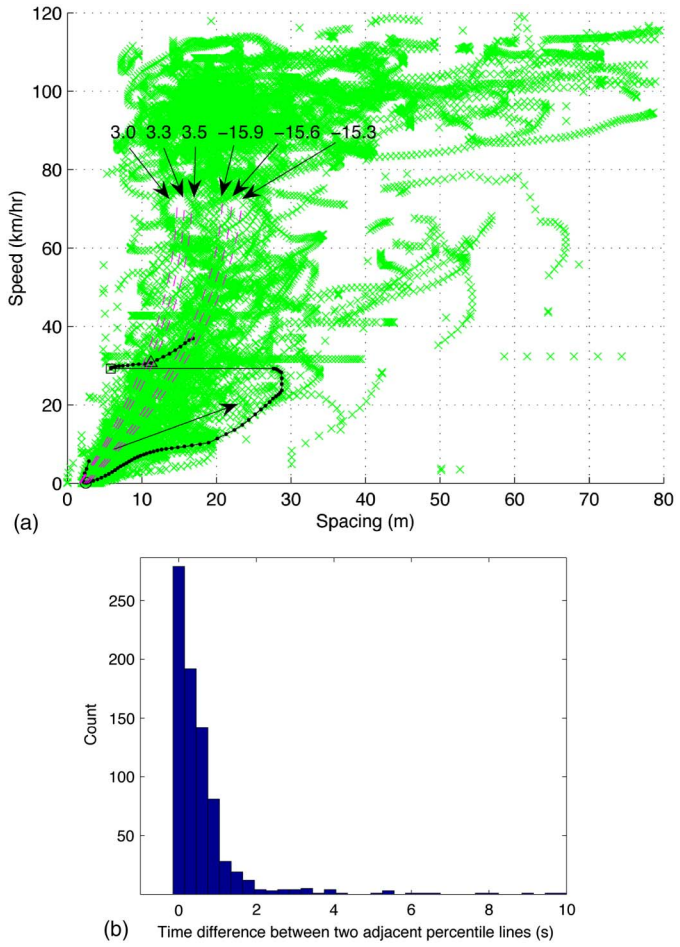

Fig. 8. Lane-change accommodation example: (a) speed–spacing relationship as seen on one run, highlighting the portion surrounding an LCM; the six numbers show the times (in seconds) when the highlighted curve crosses the respective percentile lines; the times are relative to the actual LCM time; (b) Ddistribution from all LCM of the time difference between adjacent percentile lines

The cloud of points in Fig. 8(a) shows all of the speed–spacing measurements recorded on a travel-time run from November 16, 2005. The dark set of points highlights the time-series progression from this cloud immediately before and after an LCM, while the dashed lines show the percentile curves from Fig. 7(d). The arrow shows the progression of the highlighted time-series data and there is approximately 0.3 s between each data point in the series due to the LIDAR sampling rate. The square denotes the instant of the LCM, in this case a vehicle enters the lane, and the spacing is reduced after the LCM. As illustrated in this example, the following driver may begin accommodating by creating a gap and taking a longer spacing before the entrance is recorded. Whenever this situation occurs, at the instant the entrance is recorded, the spacing abruptly jumps from the right of the quasi-equilibrium state to the left of it. If such a jump occurs it is not taken as the end of the lane-change accommodation because the time-series does not return to the quasi-equilibrium states at this instant. The six numbers above the percentile curves indicate the given crossing time in seconds relative to the instant the LCM occurred, i.e., the instant when the time series crossed the respective curve. The crossing times of the 30th, 35th, and 40th percentile curves are all quite close (3.0, 3.3, and 3.5 s, respectively), similarly, the crossing times of the 60th, 65th, and 70th percentile curves are also quite close (−15.9, −15.6, and −15.3 s, respectively), indicating that the choice of the specific percentile thresholds is not critical in this case.

Drivers may not make accommodations for an LCM ahead of them when traveling at free-flow speed because the lead vehicle is so far away that there is no interaction between the vehicles. To ensure that drivers are car-following, this research is limited to congested conditions (below 72 km/h). Using all of the complete observed LCMs when the speed is below 72 km/h, Fig. 8(b) shows the distribution of the difference of crossing times between successive percentile curves (30th to 35th, 35th to 40th, 60th to 65th, and 65th to 70th). The median time to cross the successive percentile curves is 0.3 s, and 86% of the successive crossing times are within 1 s. As noted earlier in the section on identifying lane-change maneuvers, the present work excludes any LCM that does not directly involve the probe vehicle. Out of the 167 LCMs with speed less than 72 km/h from all data sets, only 61 LCMs have a complete accommodation process. The remaining LCMs are excluded from further analysis due to the following reasons: 86 LCMs are interrupted by another LCM; 26 LCMs transition into a free-flow state before the end of the LCAT; finally 45 LCMs the time-series speed–spacing relation itself is not complete because the LIDAR lost the lead vehicle and, thus, no spacing is available.

Results and Analysis

In the case of LCMs immediately in front of the probe vehicle, this research differentiates between when another vehicle enters the probe vehicle's lane (entering vehicle) and when the lead vehicle departs probe vehicle's lane (departing vehicle). The Fig. 9(a) shows the CDF of LCAT for entering and departing vehicles. The CDF of LCAT for departing vehicles is predominantly to the right of the CDF for entering vehicles, i.e., the LCAT for departing vehicles is typically larger than that for entering vehicles.

J. Transp. Eng. 2012.138:1051-1061.

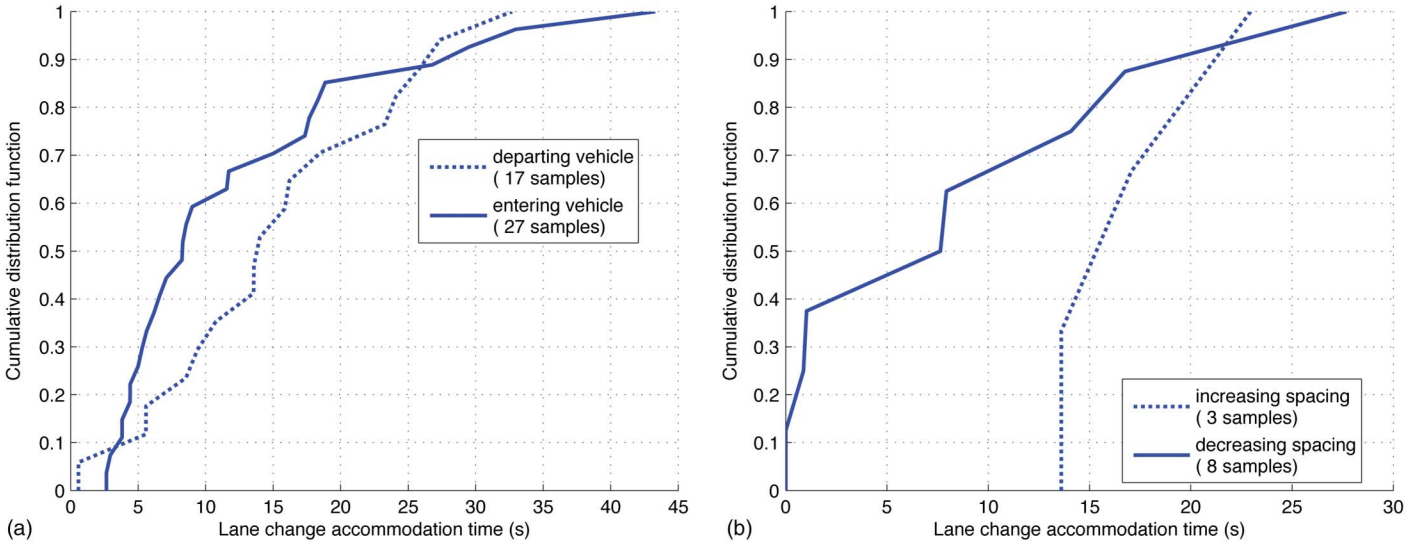

Fig. 9. CDF of accommodation time: (a) LCMs ahead of the probe vehicle; (b) LCMs by the probe vehicle

Similarly, when the probe vehicle undertakes an LCM, this research differentiates between whether the lead vehicle in the new lane is closer (decreasing spacing) or farther (increasing spacing) than the lead vehicle in the old lane. Fig. 9(b) shows the CDF of LCAT for decreasing and increasing spacing maneuvers; here the CDF of LCAT for increasing spacing is predominantly to the right of the CDF for decreasing spacing, i.e., the LCAT for increasing spacing is typically larger than that for decreasing spacing.

Both plots in Fig. 9 exhibit an asymmetry between the respective pair of CDFs. Tables 1 and 2 show the mean and median of the LCAT for the distributions from Fig. 9, for those LCMs ahead of the probe vehicle and by the probe vehicle, respectively. In either case the results are intuitive; in the case of an entering vehicle or decreasing spacing, the probe vehicle driver has no choice but to respond quickly or risk following the new lead vehicle at an unsafe spacing. But in the case of a departing vehicle or increasing spacing, the driver can safely take his or her time to consume the excess spacing and return to the quasi-equilibrium state. This asymmetry is consistent with Wang and Coifman (2008), reaffirming their findings that this imbalance is one source for disturbance formation and growth in queues. The LCAT in the present work are in general higher than those in Wang and Coifman, underscoring the fact that LCAT is highly dependent on many parameters, and the use herein is only meant for relative comparisons when the parameters are held constant.

Table 1. Summary Statistics from the Distributions of the Lane-Change Accommodation Process First Shown in Fig. 9(a) for LCM ahead of the Probe Vehicle

LCAT	Mean	Median	Standard deviation
Departing vehicle (17 samples)	15.6	14.0	8.8
Entering vehicle (27 samples)	12.4	8.3	10.3

Note: All times are in seconds.

Table 2. Summary Statistics from the Distributions of the Lane-Change Accommodation Process First Shown in Fig. 9(b) for LCM by the Probe Vehicle

LCAT	Mean	Median	Standard deviation
Increasing spacing (3 samples)	17.9	17.0	4.7
Decreasing spacing (8 samples)	9.5	7.8	9.6

Note: All times are in seconds.

Conclusions

Lane-change maneuvers (LCMs) have been suspected of being a source of traffic disturbances. To date there has been limited research on the microscopic impacts of LCMs on traffic flow due to the difficulty in collecting the necessary data. This paper employs an instrumented probe vehicle to extend the microscopic analysis beyond the limited periods and spatial coverage available from the few publicly available microscopic vehicle trajectory data sets, namely those from NGSIM and Turner Fairbanks.

The first objective of this paper is to provide independent validation of the findings reported by Wang and Coifman (2008). Wang and Coifman found that the LCAT was imbalanced between the exited lane and the entered lane; thus, an LCM within a queue induces a ripple in the traffic state that propagates upstream. As summarized above, Wang and Coifman described how this imbalance is one source for disturbances to form and grow in queues, potentially being a source of stop-and-go traffic. But Wang and Coifman only used data from one hour, in one lane, on one facility. The present study used approximately 90 h of data, along an extended corridor, but limited to a small number of vehicles around the probe vehicle. Within this set, there were 167 LCMs during congestion, but only 61 were uninterrupted. The new means to measure the LCAT presented in this paper yields results consistent with Wang and Coifman. Ultimately the empirical results from these studies should help develop more robust microscopic LCM models that better capture the impacts of LCMs on traffic.

The methodology of extracting information from the probe vehicle data is just as important as the specific traffic phenomena observed, because these tools will be of value in other studies. The process of generating a reference trajectory to provide a common reference frame to many runs through a corridor should benefit various floating car studies and potentially even emerging cell phone tracking or other active probe vehicle data streams. Though if the drivers are free to choose any lane, the methodology for generating the reference trajectory may need to be modified to use the modes of lateral displacement and look for an integer number lane

J. Transp. Eng. 2012.138:1051-1061.

widths. Next, the process of identifying both MLC and DLC by the probe vehicle to establish the lateral position across the roadway should prove to be equally beneficial. Once the probe vehicle's LCMs have been accounted for, the process of identifying LCMs from the surrounding vehicles in the LIDAR data becomes straightforward, and this process will likely prove beneficial for similar studies in the future. Finally, the process of generating the quasi-equilibrium state among the very scattered speed–spacing data should prove beneficial for other studies as well.

Acknowledgments

The authors would like to thank the numerous individuals who contributed to this work in one form or another, from the undergraduate student drivers to the anonymous reviewers who provided valuable and constructive input to this paper.

This material is based upon work supported in part by the National Science Foundation under Grant No. 0133278.

References

Ahmed, K. I., Ben-Akiva, M., Koutsopoulos, H. N., and Mishalani, R. G. (1996). "Models of freeway lane changing and gap acceptance behavior." *Proc., 13th Int. Symp. on Transportation and Traffic Theory*, Elsevier, Oxford, UK, 501–515.

Ahn, S., and Cassidy, M. J. (2007). "Freeway traffic oscillations and vehicle lane-change maneuvers." *Proc., 17th Int. Symp. of Transportation and Traffic Theory*, Elsevier, Oxford, UK, 691–710.

BPR. (1950). *Highway capacity manual*, Bureau of Public Roads, Washington, DC.

Coifman, B., and Dhoorjaty, S. (2004). "Event data-based traffic detector validation tests." *J. Transp. Eng.*, 130(3), 313–321.

Coifman, B., Mishalani, R., Wang, C., and Krishnamurthy, S. (2006). "Impact of lane-change maneuvers on congested freeway segment delays: Pilot study." *Transportation Research Record 1965*, Transportation Research Board, Washington, DC.

FHWA. (2006a). "Interstate 80 freeway dataset." *U.S. Federal Highway Administration*, *FHWA-HRT-06-137* ⟨http://www.fhwa.dot.gov/publications/research/operations/06137/index.cfm⟩ (May 24, 2012).

FHWA. (2006b). "NGSIM overview." *U.S. Federal Highway Administration*, *FHWA-HRT-06-135* ⟨http://www.fhwa.dot.gov/publications/research/operations/its/06135/index.cfm⟩ (May 24, 2012).

Gao, B., and Coifman, B. (2006). "Vehicle identification and GPS error detection from a LIDAR equipped probe vehicle." *Proc., 9th IEEE Int. Conf. on Intelligent Transportation Systems*, IEEE, Piscataway, NJ.

Gao, B., and Coifman, B. (2007). "Extracting traffic flow characteristics with an instrumented probe vehicle." *Proc., Traffic and Granular Flow 2005*, Springer, Berlin, 675–685.

Gipps, P. G. (1986). "A model for the structure of lane-changing decisions." *Trans. Res. B*, 20(5), 403–414.

Hidas, P. (2002). "Modelling lane changing and merging in microscopic traffic simulation." *Trans. Res. C*, 10(5–6), 351–371.

Hidas, P. (2005). "Modelling vehicle interactions in microscopic simulation of merging and weaving." *Trans. Res. C*, 13(1), 37–62.

Kang, K. P., and Chang, G. L. (2004). "Observations of macroscopic non-mandatory lane-changing behaviors on the Capital Beltway." *Proc., 7th Int. IEEE Conf. on Intelligent Transportation Systems*, IEEE, Piscataway, NJ, 441–446.

Laval, J. A., and Daganzo, C. F. (2006). "Lane-changing in traffic streams." *Trans. Res. B*, 40(3), 251–264.

Lighthill, M. J., and Whitham, G. B. (1955). "On kinematic waves. II. A theory of traffic flow on long crowded roads." *Proc. Roy. Soc. London, A*, 229(1178), 317–345.

Nakatsuji, T., Ranjitkar, P., and Suzuki, H. (2006). "Investigating lane-changing behavior of drivers based on test track experiments." *Proc., Transportation Research Board 85th Annual Meeting*, TRB, Washington, DC.

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

Rothery, R. W. (2001). "Car following models." *Revised Traffic Flow Theory Monograph*, Transportation Research Board. ⟨http://www.tft.pdx.edu/docs.htm⟩ (Nov. 12, 2011).

Sheu, J.-B. (1999). "A stochastic modeling approach to dynamic prediction of section-wide inter-lane and intra-lane traffic variables using point detector data." *Trans. Res. A*, 33(2), 79–100.

Sheu, J.-B., and Ritchie, S. G. (2001). "Stochastic modeling and real-time prediction of vehicular lane-changing behavior." *Trans. Res. B*, 35(7), 695–716.

Smith, S. A. (1985). "Freeway data collection for studying vehicle interactions." *U.S. Federal Highway Administration*, *FHWA/RD-85/108*, Federal Highway Administration, Reston, VA.

Smith, S. A., and Mark, E. R. (1985). "Creation of data sets to study microscopic traffic flow in freeway bottleneck sections." *Transportation Research Record 1005*, Transportation Research Board, Washington, DC.

Toledo, T., Koutsopoulos, H. N., and Ben-Akiva, M. E. (2003). "Modeling integrated lane-changing behavior." *Transportation Research Record, 1857*, Transportation Research Board, Washington, DC.

Wang, C., and Coifman, B. (2008). "The effect of lane-change maneuvers on a simplified car-following theory." *IEEE Trans. on Intelligent Transportation Systems*, 9(3), 523–535.

Yang, Q., and Koutsopoulos, H. N. (1996). "A microscopic traffic simulator for evaluation of dynamic traffic management systems." *Trans. Res. C*, 4(3), 113–129.

Zhang, Y., Owen, L. E., and Clark, J. E. (1998). "Multiregime approach for microscopic traffic simulation." *Transportation Research Record, 1644*, Transportation Research Board, Washington, DC.

J. Transp. Eng. 2012.138:1051-1061.