HCM2010

HIGHWAY CAPACITY MANUAL

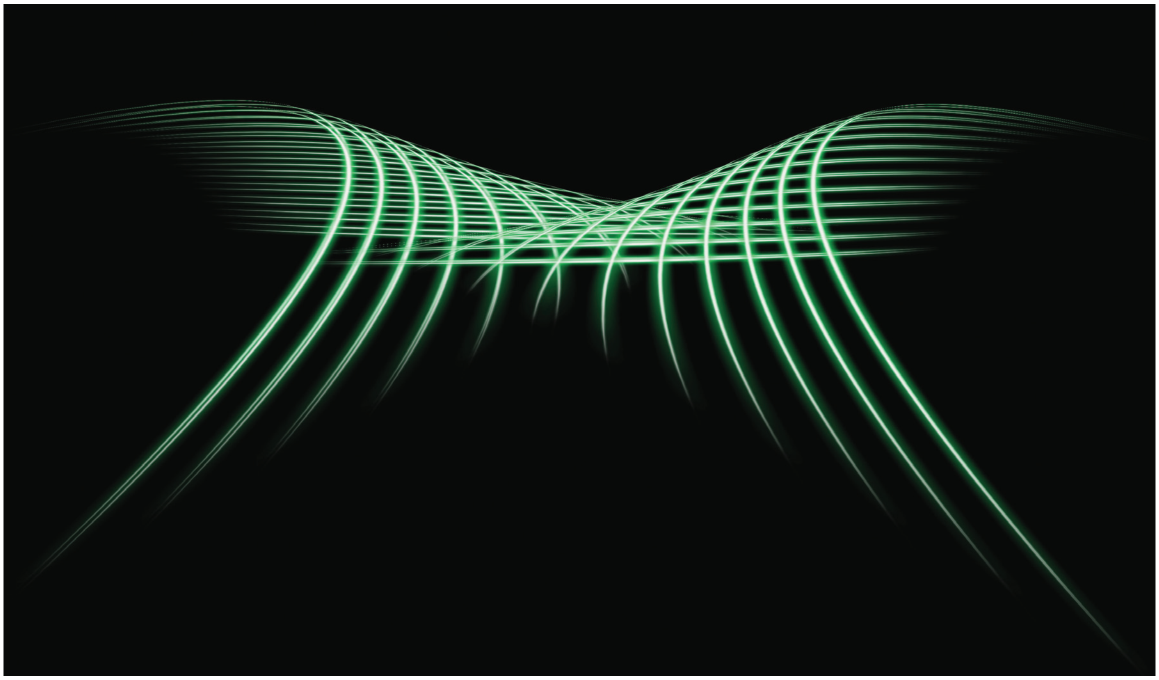

CHAPTER 24

CONCEPTS: SUPPLEMENTAL

TRB TRANSPORTATION RESEARCH BOARD
OF THE NATIONAL ACADEMIES

WASHINGTON, DC | WWW.TRB.ORG

HCM2010

HIGHWAY CAPACITY MANUAL

The HCM 2010 significantly enhances how engineers and planners assess the traffic and environmental effects of highway projects by

- Providing an integrated multimodal approach to the analysis and evaluation of urban streets from the points of view of automobile drivers, transit passengers, bicyclists, and pedestrians;
- Addressing the proper application of microsimulation analysis and the evaluation of the results;
- Examining active traffic management in relation to demand and capacity; and
- Exploring specific tools and generalized service volume tables to assist planners in quickly sizing future facilities.

The HCM 2010 consists of four volumes:

Volume 1: Concepts;
Volume 2: Uninterrupted Flow;
Volume 3: Interrupted Flow; and
Volume 4: Applications Guide (electronic only).

The four-volume format provides information at several levels of detail, to help users more easily apply and understand the concepts, methodologies, and potential applications.

Volumes 1, 2, and 3 are issued as a boxed set. Volume 4 is electronic only, accessible to registered HCM users via the Internet, and includes four types of content: supplemental chapters on methodological details and emerging issues; interpretations, clarifications, and corrections; comprehensive case studies; and a technical reference library.

To order HCM 2010, go to **http://books.trbbookstore.org/hcm10.aspx.**

For more information about HCM 2010 and TRB publications, contact the Transportation Research Board Business Office, 500 Fifth Street, NW, Washington, DC 20001 (telephone 202-334-3213; fax 202-334-2519; e-mail TRBsales@nas.edu; or through the Internet, www.trb.org).

CHAPTER 24
CONCEPTS: SUPPLEMENTAL

CONTENTS

LIST OF EXHIBITS

1. VEHICLE TRAJECTORY ANALYSIS

INTRODUCTION

Overview

This chapter of the *Highway Capacity Manual* (HCM) contains expanded guidance for the use of alternative traffic analysis tools (mostly microsimulation tools) in assessing the performance of highway facilities. An important part of the guidance deals with the use of vehicle trajectory analysis as the "lowest common denominator" for comparing performance measures from different tools. Material on vehicle trajectory analysis is also included in the following chapters:

- *Chapter 4, Traffic Flow and Capacity Concepts*, introduces the concept of individual vehicle trajectory analysis. A growing school of thought suggests that comparing results between traffic analysis tools and methods is possible only through analyzing vehicle trajectories as the "lowest common denominator." Vehicle trajectories can be used to develop performance measures that are consistent with HCM definitions, with field measurement techniques, and with each other. Examples of vehicle trajectory plots were shown that illustrate the visual properties of vehicle trajectories.
- *Chapter 7, Interpreting HCM and Alternative Tool Results*, explores the use of vehicle trajectory analysis to define and estimate consistent performance measures. First, it introduces the mathematical properties of trajectories as an extension of the visual properties. Next, it identifies the performance measures that can be computed from trajectories and explores their compatibility with the performance measures estimated by the computational procedures presented throughout the HCM.

Chapter 7 presents general guidelines for defining and comparing measures from different traffic analysis tools. Those guidelines are expanded in this section by proposing more specific trajectory analysis procedures by which consistent performance measures can be estimated. The trajectory analysis procedures described in this section were developed and tested by postprocessing the external trajectory files produced by a typical simulation tool. The postprocessor features and the process by which the procedures were developed are described elsewhere (*1*).

Several examples of the analysis of vehicle trajectories on both interrupted- and uninterrupted-flow facilities are presented here. These examples demonstrate the complexities that can arise, for example, in multilane situations, multiphase operations, situations in which the demand exceeds the capacity, and situations in which vehicles are unable to access a desired lane because of congestion. Specific procedures are then proposed and demonstrated with additional examples.

Mathematical Properties of Vehicle Trajectories

It was pointed out in Chapter 7 that an analysis of vehicle trajectories requires a mathematical representation that includes a set of properties that are associated with each vehicle at specific points in time and space. Some of the material on mathematical properties of vehicle trajectories presented in this section was originally included in Chapter 7. It is repeated here to provide a convenient introduction to the topic of vehicle trajectory analysis. A graphic representation of the path of an individual vehicle in space and time is also repeated here as Exhibit 24-1.

Exhibit 24-1
Vehicle Data Stored for Each Time Step

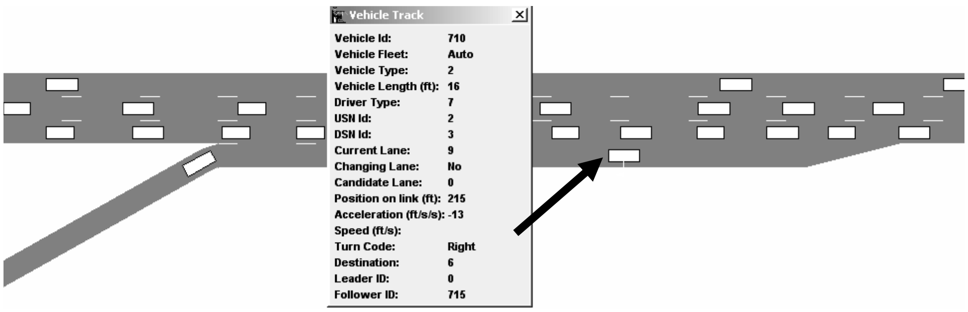

Many properties can be associated with a specific vehicle at a point in time. Some properties are required for the accurate determination of performance measures from trajectories. Others are used for different purposes, such as safety analysis.

Basic Trajectory Properties

The basic trajectory properties from which all the required performance measures can be estimated include the following information for each vehicle within the facility boundaries and for each time step within the analysis period:

- *Vehicle identification*: Vehicle identification is required to distinguish a specific vehicle from all other vehicles within the facility boundaries.
- *Position*: This property is the most basic of all, and many other properties may be derived from it. A one-dimensional position is sufficient to produce performance measures. Some question remains about a universal representation of position, because different tools specify the position in different ways. A common reference point for position needs to be established. A reference point that indicates the relative position of the vehicle in the link would be desirable to enable developers to produce uniform measures.
- *Link or segment*: A link or segment is required to associate performance measures with a specific link or analysis segment for reporting purposes.
- *Lane*: In multilane facilities, it is important to know the lane in which the vehicle is traveling, because headways, densities, and other measures must be estimated by lane. It is also necessary for identifying lane changes.

Static Vehicle and Facility Parameters

Some required properties can be derived from the basic properties with knowledge of certain parameters that are constant with respect to time, including the following:

- *Vehicle length*: Required to convert headways to gaps, and
- *Link end positions*: Required to determine the position of the vehicle with respect to the upstream or downstream end of the link.

Some simulation tools repeat this static information in each record to avoid the need for an external parameter file.

Derived Trajectory Properties

The remainder of the required trajectory properties can be derived from the basic properties as follows:

- *Instantaneous speed*: This property can be determined from the relative positions of the vehicle at time t and time $t - \Delta t$, assuming a constant acceleration during Δt. However, since most tools update vehicle positions from the speeds, it is common to include speed as a basic trajectory property.
- *Instantaneous acceleration*: This property can be determined from the relative speeds of the vehicle at time t and time $t - \Delta t$, assuming a constant acceleration during Δt. However, since most tools update vehicle speeds from the acceleration, it is common to include acceleration as a basic trajectory property.

VEHICLE TRAJECTORY ANALYSIS MECHANISM

A common analysis mechanism is required to produce performance measures from vehicle trajectory files. A software utility, Vehicle Trajectory Analysis for Performance Evaluation (VTAPE), was developed by National Cooperative Highway Research Program (NCHRP) Project 3-85 for this purpose. The operation of this utility program is described elsewhere (*1*).

VTAPE provides detailed step-by-step analyses with intermediate values reported for all parameters at each time step. It is a useful tool for understanding and enhancing the trajectory analysis methodology. However, it is also intended as a research tool and should not be viewed as an end-user analysis tool. Its ultimate purpose will be to serve as the computational engine for the continuing analysis and development of vehicle trajectory procedures. It was developed with that purpose in mind.

Vehicle Trajectory Data Source

Simulation tools propagate vehicles through a roadway segment by keeping track of the trajectory properties that are maintained internally within the traffic flow model. Many simulation tools produce files containing vehicle trajectory data in some form for their own internal use. The trajectory files produced by a commonly used simulation tool were used to create the examples presented in this chapter. The trajectory file records the location of each vehicle in a single simulation run for every time step of the simulation in a published format. The files are used internally by the simulation tool to support animated graphics displays. All the data required to produce performance measures, as identified in Exhibit 24-1, are included in the trajectory files.

VTAPE Operation

A full description of VTAPE and instructions for its installation and operation are provided elsewhere (*1*). The capabilities of this software as they apply to the examples in this section are summarized below.

With VTAPE, it is possible to select up to eight links forming a continuous route for analysis. The data for the selected links are read from the animated graphics files. The selected links are joined together for plotting and analysis to form a linear route. Each link is configured in terms of the following:

- Upstream and downstream nodes that define the link,
- Link length,
- Free-flow speed, and
- Control at the downstream end of the link.

VTAPE also has capacity for up to three auxiliary lanes per link to accommodate acceleration or deceleration lanes on freeways and turn bays on surface streets. The auxiliary lane properties include the following:

- Lane number assigned by the simulation tool,
- Auxiliary lane length,
- Free-flow speed, and
- Type of lane (left turn, right turn, acceleration, or deceleration).

VTAPE performs the following functions, all of which are demonstrated in this section:

1. Trajectory plotting, either by lane or for all lanes in the link.
2. Longitudinal analysis of the trajectory of vehicles as they traverse a link. A single vehicle may be chosen for a detailed analysis to illustrate the analysis procedure, or all vehicles traversing the link in a given time period may be analyzed separately, with their performance measures included. The measures determined by this type of analysis include various types of delay measures and stop-related measures.
3. Spatial analysis, which involves consideration of all the vehicles on a link at a specific time step. The two principal spatial measures are density and queue lengths.

TRAJECTORY ANALYSIS EXAMPLES

As a prelude to development of the actual procedures for producing the performance measures presented in Chapter 7, several examples of vehicle trajectory analysis with VTAPE were prepared to demonstrate the ability to quantify the trajectory properties. These examples are now presented and discussed separately.

Basic Signalized Intersection

The first example is very basic. The intersection configuration involves two single-lane, one-way streets as shown in Exhibit 24-2. To simplify the situation even more, the simulation parameters are adjusted to enforce a uniform operation. Essentially, all the randomness inherent in simulation is removed. A simulation of uniform conditions would not normally produce useful results, but this example provides a good starting point to illustrate the nature of vehicle trajectory plots.

Exhibit 24-2
Basic Signalized Intersection Example

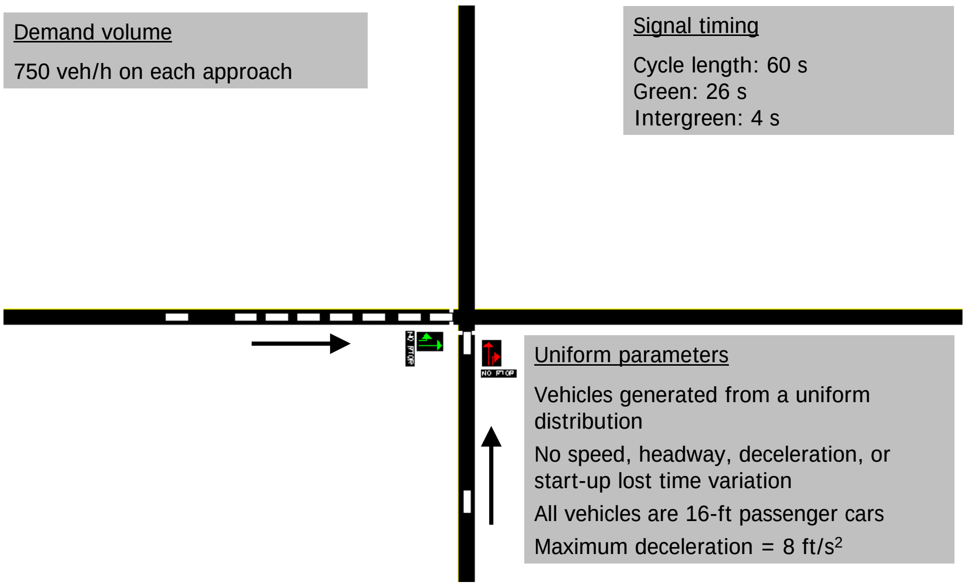

A trajectory plot showing two cycles of simulated operation for this example is presented in Exhibit 24-3(a). This form is the classic one that appears often in the literature to support discussion related to queue accumulation and discharge. A copy of the exhibit used in Chapter 31, Signalized Intersections: Supplemental, to illustrate the basic traffic signal principles is also included as Exhibit 24-3(b). The difference between the two figures is that one was produced directly from the vehicle trajectory data and the other was drawn by hand. The ability to reproduce the classic representation from controlled conditions will provide a measure of confidence in the validity of future examples involving much more complicated situations.

Restoring Randomness to the Simulation

To simplify the discussion, the first example was presented with all randomness removed from the operation. Subsequent examples are more realistic in their treatment of traffic flow. Vehicles are generated at entry points from a Poisson distribution, and the simulation tool's default parameters for randomizing driver behavior are applied.

Exhibit 24-3
Trajectory Plots for Uniform Arrivals and Departures

Note the similarity between the trajectories obtained from the file (above) and those that were developed manually in Chapter 31 (below) to illustrate the basic principles of signalized intersection operation.

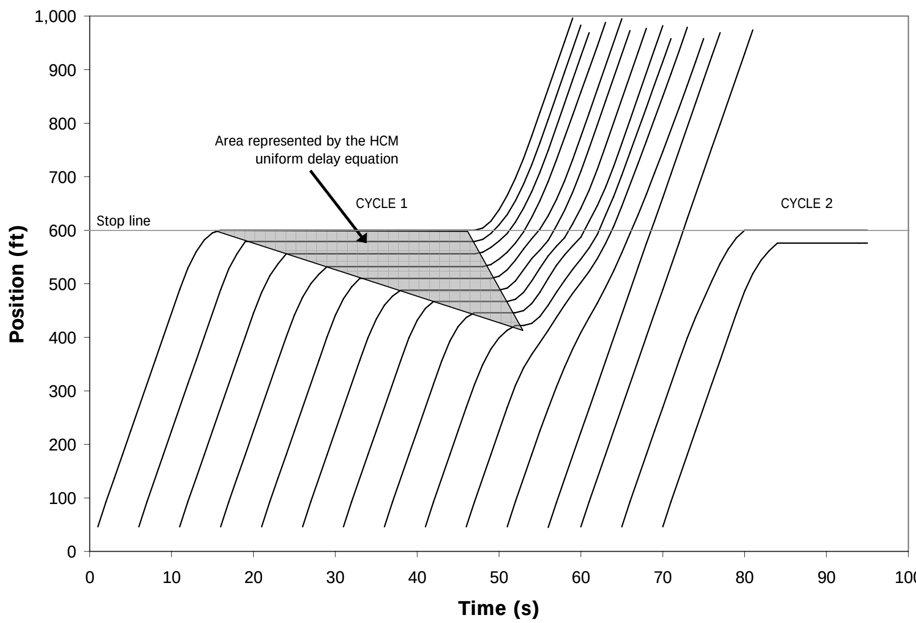

(a) Plot Produced from Simulation

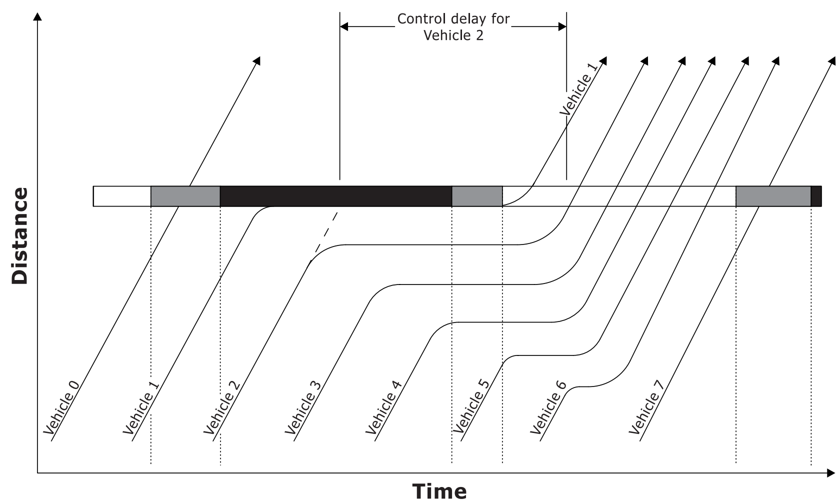

(b) Plot Produced by Hand

Exhibit 24-4 shows a sample trajectory plot for the same operation depicted in Exhibit 24-3. As expected, the individual trajectories follow the same pattern as the uniform case, except that some spacings and speeds are not as consistent. The trajectory lines do not cross each other in this example because the example uses a single-lane approach and overtaking is not possible.

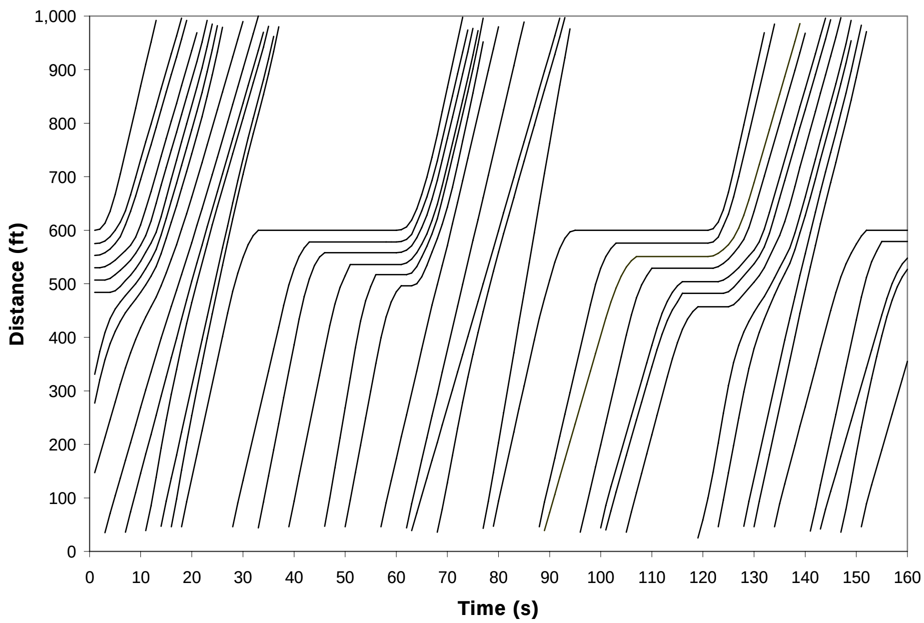

Exhibit 24-4
Introducing Randomness into the Simulation

Vehicle Trajectories for Oversaturated Operation

Up to this point, the examples have involved volume/capacity (*v/c*) ratios less than 1.0, in which all vehicles arriving on a given cycle were able to clear on the same cycle. Saturation levels close to and above 1.0 present a different picture. Three cases are presented here:

1. *Cycle failure,* occurring when saturation approaches 1.0 and residual queues build on one cycle but are resolved on the next cycle;
2. *Oversaturated operation,* a situation in which the link has a demand volume exceeding the link's capacity and queues extend throughout the approach link; and
3. *Undersaturated operation,* in which queues extend to an upstream link for a part of a cycle because of closely spaced intersections.

Cycle Failure

A cycle failure example is presented in Exhibit 24-5. This trajectory plot shows a situation in which some vehicles arriving in Cycle 1 were unable to clear until Cycle 2. This condition is identified from the trajectory plot for four stopped vehicles (i.e., horizontal trajectory lines) that were forced to stop again before reaching the stop line. These vehicles became the first four vehicles in the queue for Cycle 2. Fortunately, the arrivals during Cycle 2 were few enough that all stopped vehicles were able to clear the intersection before the beginning of the red phase. A closer inspection of Exhibit 24-5 shows that one more vehicle, which was not stopped, was also able to clear.

Exhibit 24-5
Cycle Failure Example

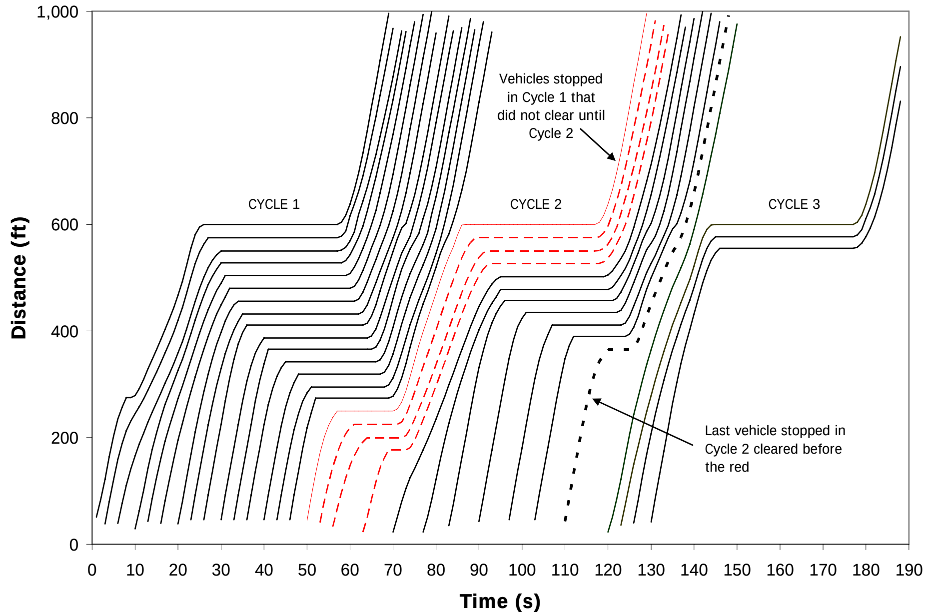

Severely Oversaturated Operation

Oversaturated operation was produced by increasing the demand volume to the point where it exceeded the capacity of the approach. The increased demand produced a queue that extended the length of the link. Inspection of the animated graphics showed that the queue did, in fact, back up beyond the link entry point.

The vehicle trajectory plot for this operation is presented in Exhibit 24-6. The move-up process is represented in the trajectories. Vehicles entering the link require up to three cycles to clear the intersection. The implications on control delay computations when the queue occupies a substantial proportion of the link are discussed in Chapter 7.

Exhibit 24-6
Oversaturated Signal Approach

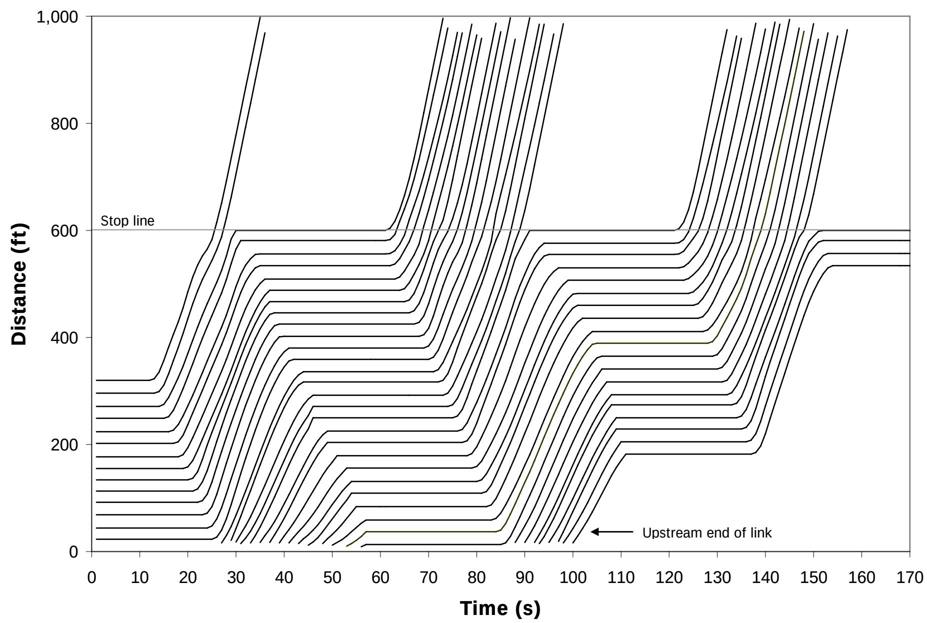

A larger question is what to do with the vehicles denied entry during the analysis period. The answer is that, as indicated in Chapter 18, Signalized Intersections, the analysis period must be long enough to include a period of uncongested operation at each end. The delay to vehicles denied entry to this link will be accounted for in upstream links during the period. The upstream links must include a holding area outside the system. Some tools include the delay to vehicles denied entry and some do not. If a tool is used that does not include denied-entry delay, then fictitious links must be built into the network structure for that purpose.

Queue Backup from a Downstream Signal

Even when an approach is not fully saturated, queues might back up from a downstream signal for a portion of the cycle. It happens when intersections are closely spaced. An example of queue backup within a cycle is shown in Exhibit 24-7.

The two-intersection configuration for this example is shown in Exhibit 24-7(a). The graphics screen capture shows that vehicles that would normally pass through the upstream link are prevented from doing so by queues that extend beyond the end of the downstream link for a portion of the cycle. The question is how to treat the resulting delay.

Exhibit 24-7
Queue Backup from a Downstream Signal

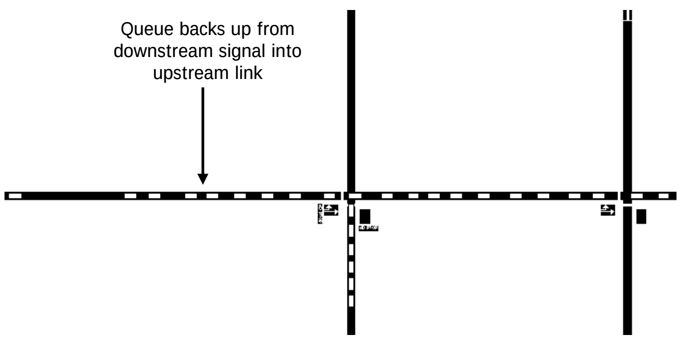

(a) Simulation Graphics Representation

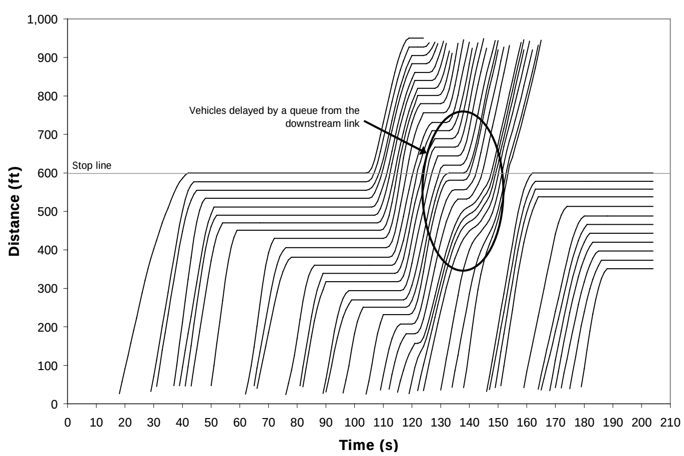

(b) Vehicle Trajectory Representation

By the definitions given to this point, the delay in the upstream link would be assigned to the upstream link, even though the signal on the downstream link was the primary cause. The important thing is not to overlook any delay and to assign all delay somewhere and in a consistent manner. With simulation modeling, the only practical place to assign delay consistently is the link on which the delay occurred. Subtle complexities make it impractical to do otherwise. For example, the root cause of a specific backup might not be the immediate downstream link. In fact, the backup might be secondary to a problem at some distant location in the network at some other point in time.

More Complex Signal Phasing

Up to this point only simple signal phasing has been considered. Many applications involve simulating more complex phasing on urban streets. As an example of a more complex situation, a left turn moving on both a protected and a permitted phase is examined.

Exhibit 24-8 shows the trajectory plot for an eastbound left-turn movement from an exclusive lane, controlled by a signal with both protected and permitted phases. In this case, the upstream link is the eastbound approach to the intersection and the downstream link is the northbound approach to the next intersection. Because the distance on a trajectory plot is one dimensional, the distance scale is linear, even though the actual route takes a right-angle bend.

Exhibit 24-8
Trajectory Plot for More Complex Signal Phasing

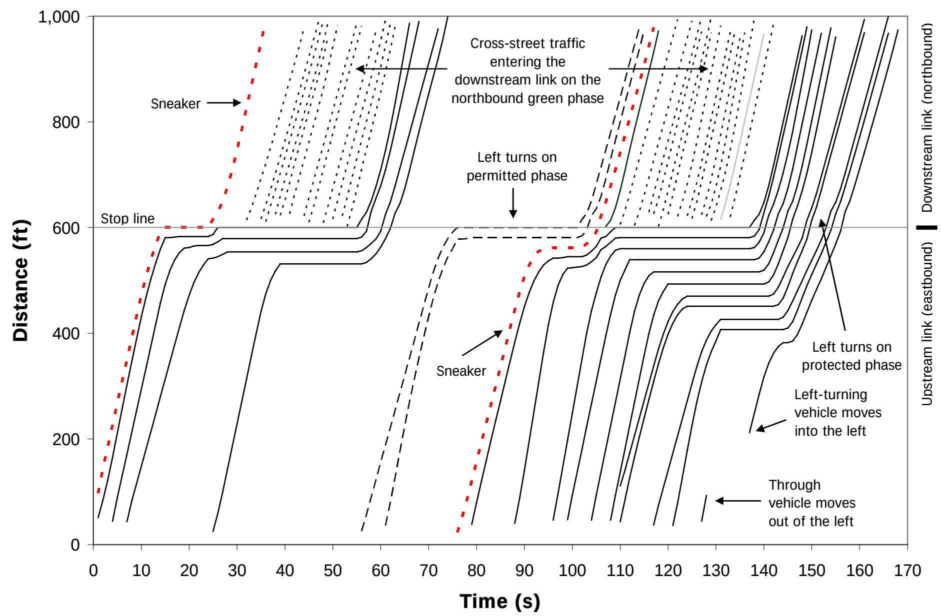

Even with an undersaturated operation, this trajectory plot is substantially more involved than the previous ones. Several phenomena are identified in the figure, including the following:

1. Cross-street traffic entering the downstream link on the northbound phase: These vehicles do not appear on the upstream link because they are on a different link. They enter the downstream link at the stop line on the red phase for the left-turn movement of interest.

2. Left turns on the protected phase, shown as solid lines on the trajectory plot: The protected left-turn phase takes place immediately after the red phase. The left-turning vehicles begin to cross the stop line at that point.
3. Left turns on the permitted phase, shown as broken lines on the trajectory plot: The permitted left-turn phase takes place immediately after the protected phase. There is a gap in the trajectory plot because the left-turning vehicles must wait for oncoming traffic to clear.
4. Left-turn "sneakers": It is not possible to identify a sneaker explicitly on the trajectory plot; however, the last left turn to clear the intersection on the permitted phase is probably a sneaker if it enters at the end of the permitted phase.
5. Left-turn vehicles that enter the link in the through lane and change into the left lane somewhere along the link: These vehicles are identified by trajectories that begin in the middle of the link.
6. Through vehicles that enter the link in the left-turn lane and change into the through lane somewhere along the link: These vehicles are identified by trajectories that end abruptly in the middle of the link.

The trajectory plot shown for this example is more complex than the previous plots; however, performance can be analyzed in the same way.

Freeway Examples

Freeway trajectories follow the same definitions as surface street trajectories, but the queuing patterns differ because they are created by car-following phenomena and not by traffic signals. The performance measures of interest also differ. There is no notion of control delay on freeways because there is no control. The level of service on uninterrupted-flow facilities is based on traffic density expressed in units of vehicles per mile per lane. In some cases, such as merging segments, the density in specific lanes is of interest.

Two cases are examined. The first case deals with a weaving segment, and the second case deals with merging at an entrance ramp.

Weaving Segment Example

Simulation Network Structure

The problem description, link–node structure, and animated graphics view for the weaving segment example are shown in Exhibit 24-9. It is the same scenario used in Example Problem 1 in Chapter 12, Freeway Weaving Segments. There are two lanes on the freeway and on each ramp. The two ramp lanes are connected by full auxiliary lanes.

Exhibit 24-9
Weaving Segment Description and Animated Graphics View

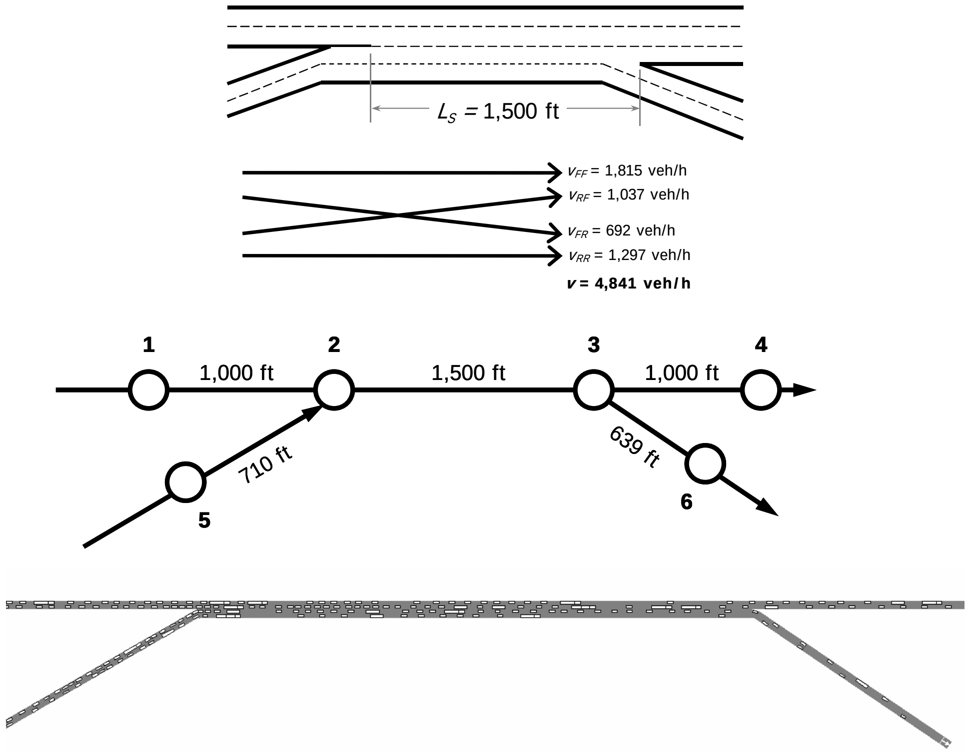

Note: L_S = length of segment, V_{FF} = vehicles entering from freeway and leaving to freeway, V_{RF} = vehicles entering from ramp and leaving to freeway, V_{FR} = vehicles entering from freeway and leaving to ramp, V_{RR} = vehicles entering from ramp and leaving to ramp, veh/h = vehicles per hour.

Vehicle Trajectories for the Freeway Lanes

The vertical (i.e., distance) axis of the trajectory plot provides a linear one-dimensional representation of a series of connected links. The links can follow any pattern as long as some of the vehicles leaving one link flow into the next link. The VTAPE utility accommodates a maximum of eight connected links. When multiple links are connected to a node (as is usually the case), then different combinations of links may be used to construct a multilink trajectory analysis. The route configuration must be designed with the end product in mind. Sometimes it is necessary to examine multiple routes to obtain a complete picture of the operation.

There are two entry links and two exit links to the weaving segment, giving four possible routes for analysis. Two routes are examined in this example. The first route, which is represented in Exhibit 24-10, shows the traffic entering the weaving segment from the freeway and leaving to the freeway (V_{FF} in Exhibit 24-9), represented by Links 1–2–3–4.

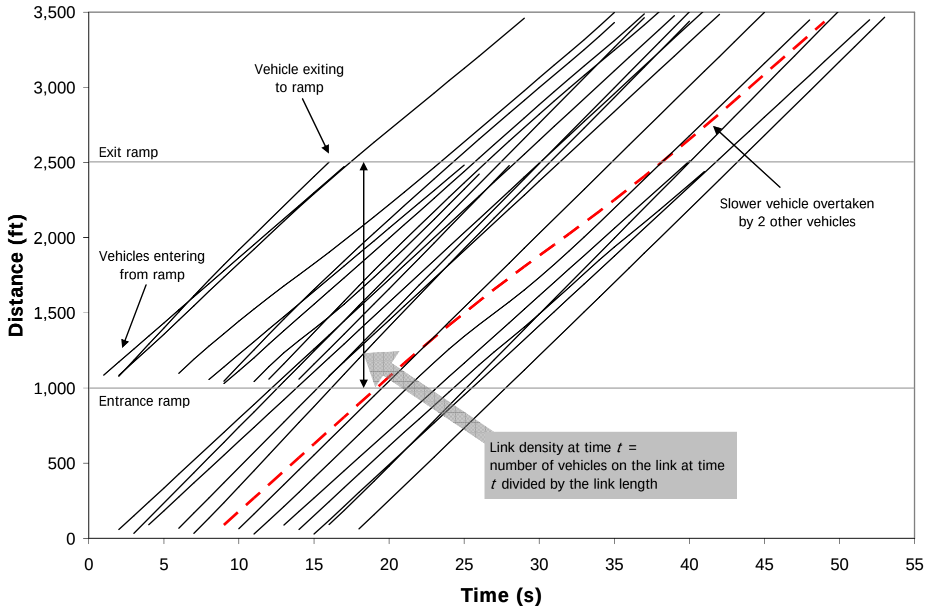

Exhibit 24-10
Trajectory Plot for Freeway Links

This plot is a multilane plot so, unlike previous plots, some of the trajectory lines might cross each other because of different speeds in different lanes. One such instance is highlighted in Exhibit 24-10. This figure also shows vehicles that enter and leave the weaving segment on the ramps. Because the ramps are not part of the selected route, the ramp vehicles appear on the trajectory plot only on the link that represents the weaving segment. Examples of ramp vehicles are identified in the figure.

The definition of link density (vehicles per mile) is also indicated in Exhibit 24-10. Density as a function of time *t* is expressed in vehicles per mile and is determined by counting the number of vehicles within the link and dividing by the link length in miles. Average lane density (vehicles per mile per lane) on the link may then be determined by dividing the link density by the number of lanes. To obtain individual lane densities, it is necessary to perform the trajectory analysis on each lane. The analysis must also be performed on a per lane basis to examine individual vehicle headways.

Vehicle Trajectories for the Entrance and Exit Ramps

By specifying the links on the route as 5–2–3–6 instead of 1–2–3–4, the trajectories for vehicles entering and leaving the weaving segment on the ramps (V_{RR} in Exhibit 24-9) can be examined. This trajectory plot is shown in Exhibit 24-11. This figure is very similar to Exhibit 24-10, except that the vehicles that do not appear outside the weaving segment are those on the freeway links instead of the ramp links.

It is also possible to construct two other routes, one for vehicles entering from the freeway and leaving to the exit ramp, V_{FR}, as 1–2–3–6, and for those entering from the ramp and leaving to the freeway, V_{RF}, as 5–2–3–4. These plots are not included here.

Exhibit 24-11
Trajectory Plot for Entrance and Exit Ramp Links

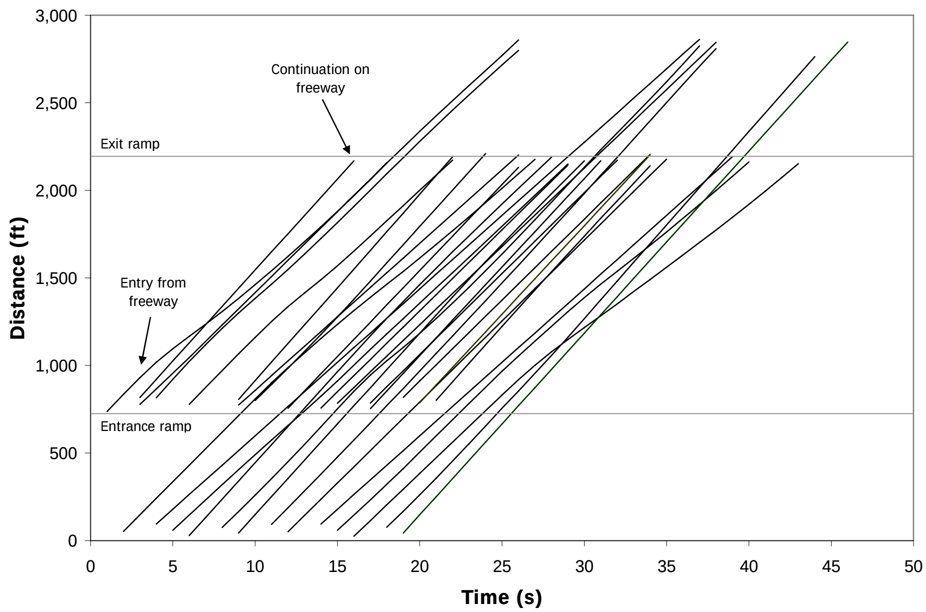

Entrance Ramp Merging Example

Merging segments provide another good example of vehicle trajectory analysis on a freeway. The merging vehicles affect the freeway operation differently in each lane, so it is necessary to examine each lane independently.

Simulation Network Structure

The same node structure used in the weaving segment example is used here. The lane configuration has been changed to be more representative of a merge operation. Three lanes have been assigned to the freeway and one lane to the entrance ramp. The demand volumes have been specified to provide a near-saturated operation to observe the effects of merging under these conditions. A graphic view of the operation is presented in Exhibit 24-12.

Exhibit 24-12
Entrance Ramp Merging Segment Graphics View

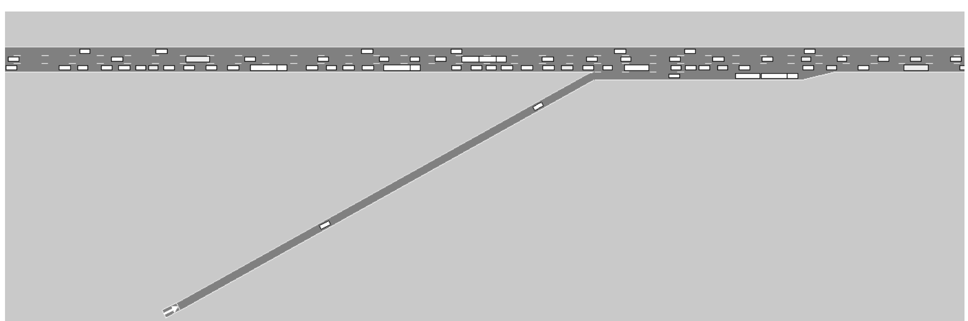

Trajectory Plots for All Lanes

Exhibit 24-13 shows the trajectory plot for all freeway lanes combined. The operation is clearly heterogeneous, with a mixture of fast and slow speeds. Many trajectory lines cross each other, and not much can be done in the way of analysis with the data from this scenario.

Exhibit 24-13
Trajectory Plot for All Freeway Lanes in the Merge Area

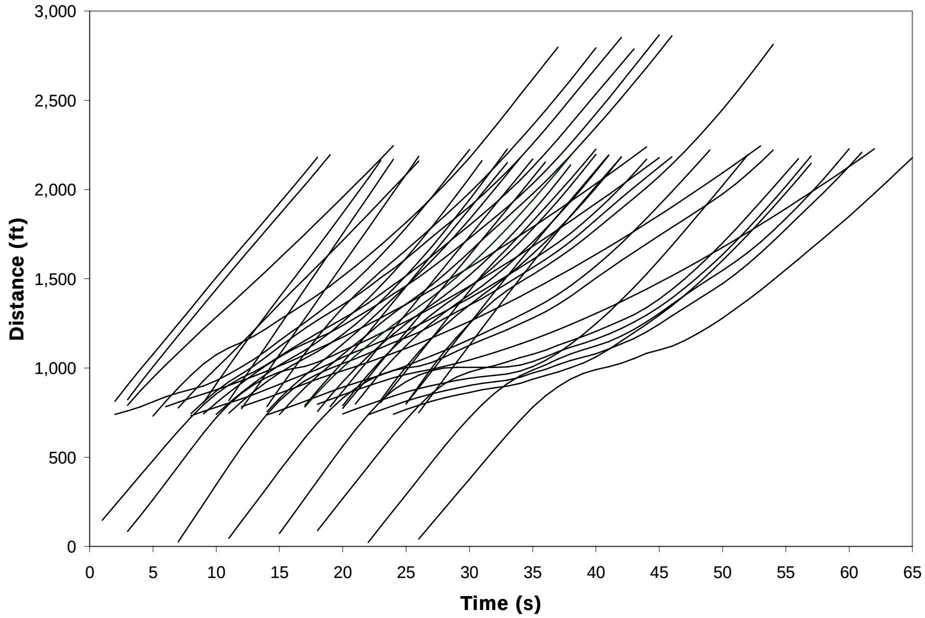

Trajectory Plots for Individual Lanes

It is clearly necessary to examine each lane individually. Exhibit 24-14, Exhibit 24-15, and Exhibit 24-16 show the trajectory plots for Lanes 1, 2, and 3, respectively. Because these plots represent individual lanes, the trajectory lines do not cross each other. The effect of the merging operation is observable (and predictable) in these three figures.

Exhibit 24-14
Trajectory Plot for Freeway Lane 1 (Rightmost) in the Merge Area

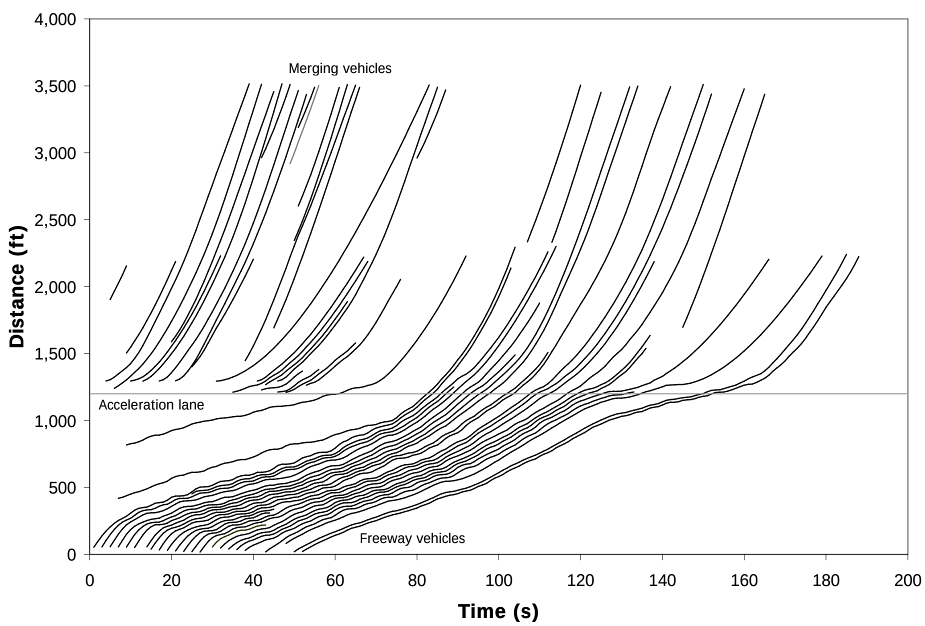

Exhibit 24-15
Trajectory Plot for Freeway Lane 2 (Center) in the Merge Area

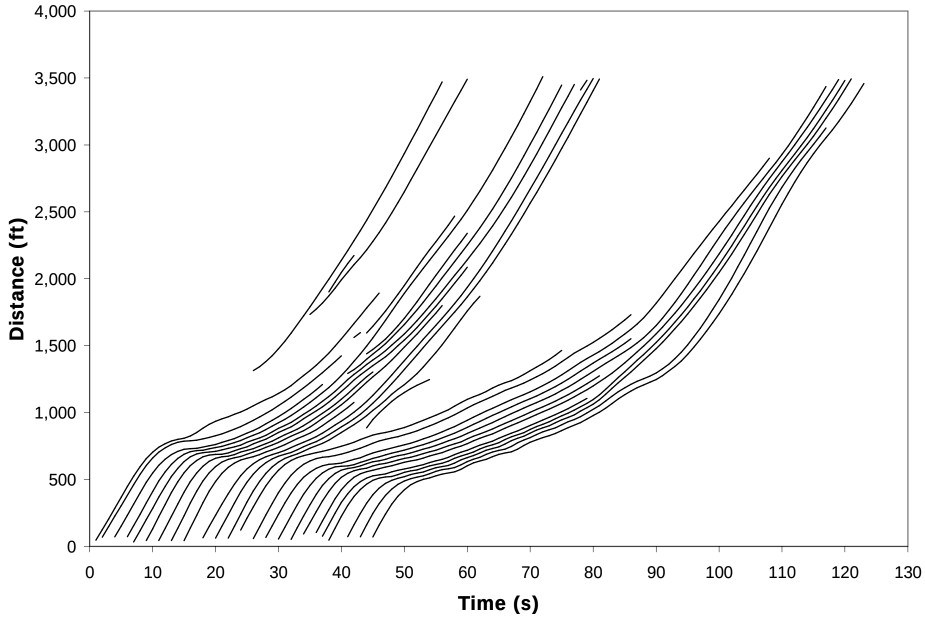

Exhibit 24-16
Trajectory Plot for Freeway Lane 3 (Leftmost) in the Merge Area

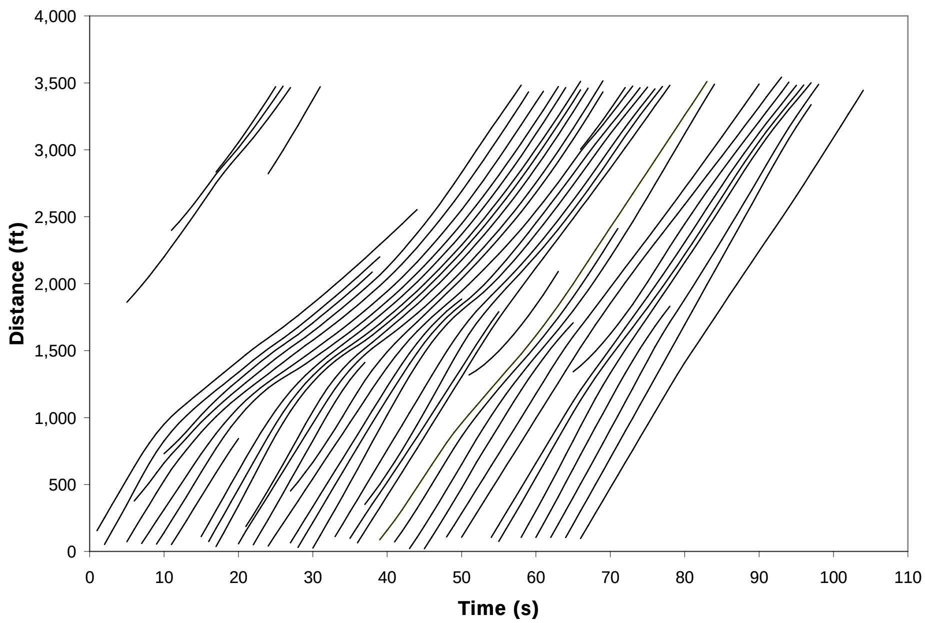

In Lane 1, freeway speeds are low upstream of the merge point. Merging vehicles enter the freeway slowly but pick up speed rapidly downstream of the merge point bottleneck. The merging vehicles enter the freeway from the acceleration lane, which begins at 1,000 ft on the distance scale. The merging vehicle trajectories before entering the freeway are not shown in Exhibit 24-14 because those vehicles are either on a different link or in a different lane.

In Lane 2, the freeway speeds are higher but still well below the free-flow speed, indicating that the merge operation affects the second lane as well. Some vehicles enter Lane 2 in the vicinity of the acceleration lane, but they are generally vehicles that have left Lane 1 to avoid the friction. Both Lane 1 and Lane 2 show several discontinuous trajectories that indicate lane changes. The

Lane 3 operation is much more homogeneous and speeds are higher, indicating a much smaller effect of the merging operation.

Trajectory Plots for Ramp Vehicles

To configure a trajectory route covering the entrance ramp vehicles, the ramp and acceleration lane, which were not represented in Exhibit 24-14 through Exhibit 24-16, must be selected in place of the upstream freeway link. The acceleration lane number must first be identified from the simulation tool's output. Because of the selected tool's unique and somewhat creative lane numbering scheme, the acceleration lane will be Lane 9. To cover both the ramp and the acceleration lane, Lane 9 must be selected on the freeway link (2–3).

The trajectory plot for this route is shown in Exhibit 24-17. The results are not quite what one might anticipate. Vehicles are observed on the ramp and in the acceleration lane, but they disappear as soon as they enter the freeway. More vehicles eventually appear toward the end of the freeway link. The vehicles disappear because Lane 9 was selected for the freeway link, so vehicles in Lane 1 do not show up on the plot. The vehicles that reappear at the end of the link are those that are leaving the freeway at the downstream exit. They reappear at that point because the deceleration lane at the end of the link is also assigned as Lane 9. This plot is not particularly useful, except that it illustrates the complexities of trajectory analysis.

Exhibit 24-17
Trajectory Plot for Acceleration and Deceleration Lanes

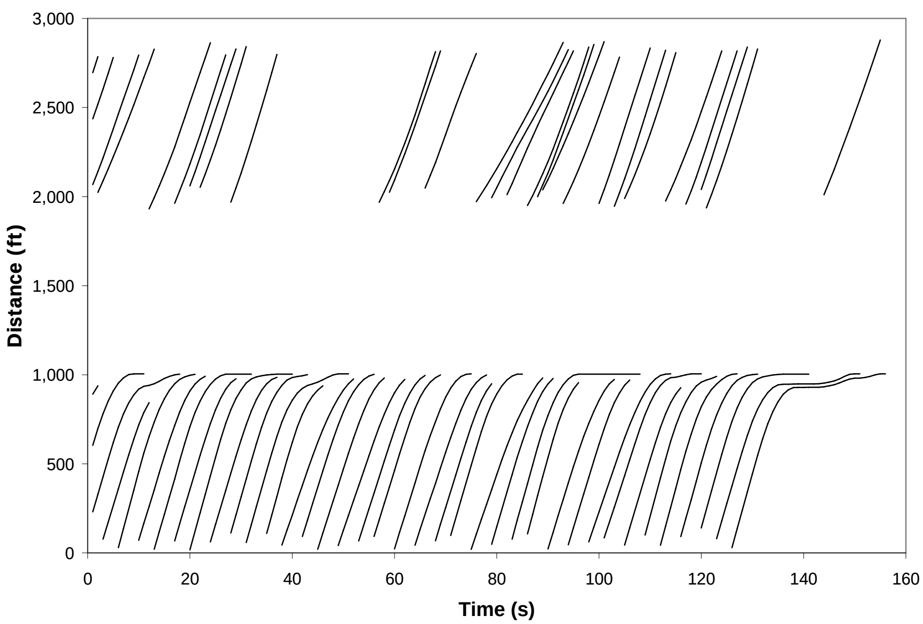

To obtain a continuous plot of ramp vehicles, nodes must be added to the network at the points where the acceleration and deceleration lanes join the freeway. These nodes are shown as Nodes 7 and 8 in Exhibit 24-18. A continuous route may then be configured as 5–2–7–8–3–4. The trajectory plot for this route is shown in Exhibit 24-19. This plot shows the entering vehicles on the ramp as they pass through the acceleration lane onto the freeway. There are some discontinuities in the trajectories because of the different point at which vehicles leave the acceleration lane.

Exhibit 24-18
Addition of Intermediate Nodes for Continuous Trajectory Plots

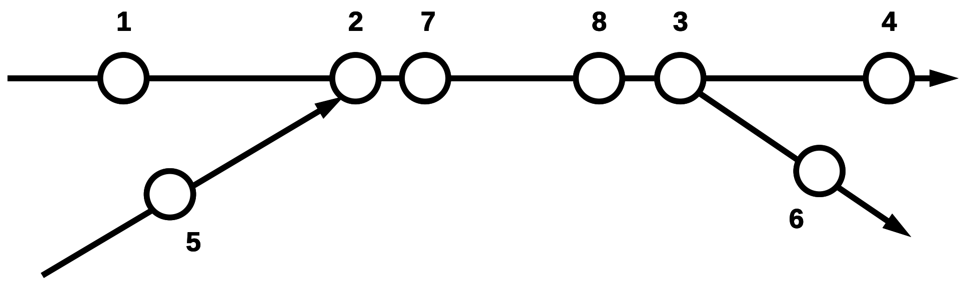

Exhibit 24-19
Trajectory Plot for Acceleration Lane and Freeway Lane 1

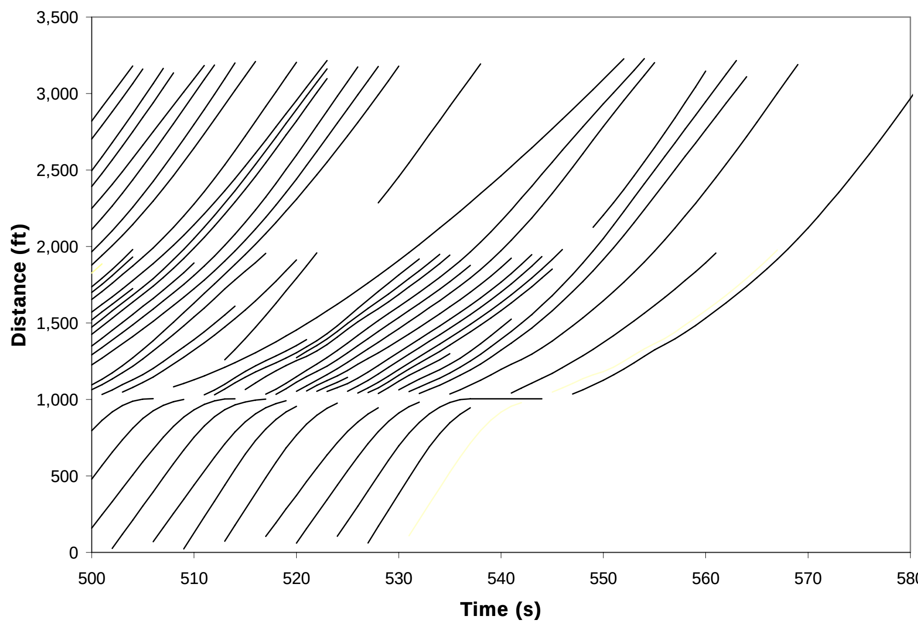

ESTIMATING PERFORMANCE MEASURES FROM VEHICLE TRAJECTORY DATA

The preceding subsections demonstrated that it is possible to produce vehicle trajectory plots that can be interpreted and analyzed. This subsection focuses more on computation of the performance measures from a mathematical analysis of the data represented in these plots.

Trajectory Analysis Procedures Overview

One development goal for this edition of the HCM was the creation of a set of computational procedures by which developers of simulation tools could produce performance measures that are consistent among different tools and, to the extent possible, compatible with the HCM's deterministic procedures. The procedures presented here were designed to be implemented easily by using the common trajectory properties described previously and illustrated by examples. Developers of simulation tools are encouraged to implement these procedures, and users of simulation tools are encouraged to consider the extent to which the procedures have been implemented in the traffic analysis tool selection process described in Chapter 6, HCM and Alternative Analysis Tools.

Requirements for Trajectory Analysis Algorithm Development

A basic set of guidelines for computing uniform performance measures from vehicle trajectory analysis was introduced in Chapter 7, Interpreting HCM and Alternative Tool Results. Since these requirements are also incorporated into the specific computational procedures proposed in this chapter, they are repeated here to promote a better understanding of the procedures. The general guidelines proposed in Chapter 7 include the following:

1. The trajectory analysis procedures are limited to analysis of trajectories produced by the traffic flow model of each simulation tool. The nature of the procedures must not suggest the need for developers to change their driver behavior or traffic flow modeling logic.
2. If the procedures for estimating a particular measure cannot be satisfactorily defined to permit a valid comparison between the HCM and other modeling approaches, then such comparisons should not be made.
3. All performance measures that accrue over time and space shall be assigned to the link and time interval in which they occur. Subtle complexities make it impractical to do otherwise. For example, the root cause of a specific delay might not be within the link or the immediate downstream link. In fact, the delay might be secondary to a problem at some distant location in the network and in a different time interval.
4. The analyst must understand that the spatial and temporal boundaries of the analysis domain must include a period that is free of congestion on all sides. This principle is also stated in Chapter 10, Freeway Facilities, and in Chapter 18, Signalized Intersections, for multiperiod signalized intersection analysis. To ensure that delays to vehicles denied entry to the system during a given period are properly recognized, it might be necessary to create fictitious links outside the physical network to hold such vehicles. A more detailed discussion of spatial and temporal boundaries is provided in Chapter 7.
5. It is important to ensure that the network has been properly initialized or "seeded" before trajectory analysis is performed. When the warm-up periods are set and applied, simulation tools typically start with an empty network and introduce vehicles until the vehicular content of the network stabilizes. Trajectory analysis should not begin until stability has been achieved. If the simulation period begins with oversaturated conditions, stability may never be achieved. See the discussion in Chapter 7 on temporal and spatial boundaries.

In addition to the general guidelines, some requirements must be addressed here to promote the development of trajectory analysis procedures that can be applied in a practical manner by the developers of simulation tools. The following requirements are suggested:

1. The algorithms must be suitable for computation "on the fly." They must not require information from a future time step that would complicate the data handling within the simulation process.

2. Arbitrary thresholds should be kept to a minimum because of the difficulty of obtaining acceptance throughout the user community for specific thresholds. When arbitrary thresholds cannot be avoided, they should be justified to the extent possible by definitions in the literature and, above all, they should be applied consistently for different types of analysis.
3. Computationally complex and time-consuming methods should be avoided to minimize the additional load on the model. Methods should be developed to simplify situations with many special cases because of the difficulty of enumerating all special cases.
4. The same definitions, thresholds, and logic should be used for determination of similar parameters in different computational algorithms for longitudinal and spatial analysis.

Summary of Computational Procedures

Several performance measures were examined in Chapter 7, and general guidelines for comparing measures produced by different tools were presented. Previous material in this section has demonstrated the potential for development of uniform measures by individual vehicle trajectory analysis and has proposed some requirements for development of the analysis procedures. Specific procedures for analyzing vehicle trajectories are now presented and demonstrated with additional examples.

Thresholds for Computation of Performance Measures

Elimination of arbitrary and user-specified values is an important element of standardization. Avoidance of arbitrary thresholds was identified earlier as a requirement for the development of trajectory analysis procedures. It is desirable to avoid all arbitrary thresholds. If thresholds cannot be avoided, they should be justified in terms of the literature. When no such justification exists, they should at least be established on the basis of consensus and applied consistently. The following thresholds cannot be avoided in vehicle trajectory analysis.

Car Length

The following is stated in Chapter 31, Signalized Intersections: Supplemental:

> *A vehicle is considered as having joined the queue when it approaches within one car length of a stopped vehicle and is itself about to stop. This definition is used because of the difficulty of keeping precise track of the moment when a vehicle comes to a stop.*

So, for estimation of queue-related measures, it is necessary to choose a value that represents one car length. For the purposes of this section, a value of 20 ft is used.

Stopped-Vehicle State

One example of an arbitrary threshold is the speed at which a vehicle is considered to have come to a stop. Several arbitrary thresholds have been applied for this purpose. To maintain consistency with the definition of the stopped state applied in other chapters of the HCM, a speed less than 5 mi/h is used here for determining when a vehicle has stopped.

Moving-Vehicle States

Other states in addition to the stopped state that must be defined consistently for vehicle trajectory analysis include the following:

- The uncongested state, in which a vehicle is moving in a traffic stream that is operating below its capacity;
- The congested state, in which the traffic stream has reached a point that is at or slightly above its capacity, but no queuing from downstream bottlenecks is present; and
- The severely constrained state, in which downstream bottlenecks have affected the operation.

These states apply primarily to uninterrupted flow. A precise definition would require very complex modeling algorithms involving capacity computations or "look ahead" features, both of which would create a computational burden. Therefore, an easily applied approximation must be sought. Threshold speeds are a good candidate for such an approximation.

It is convenient to think of these states in terms of speed ranges. To avoid specifying arbitrary speeds as absolute values, it is preferable to use the target speed of each vehicle as a reference. The target speed is the speed at which the driver prefers to travel. It differs from the free-flow speed in the sense that most simulation tools apply a "driver aggressiveness" factor to the free-flow speed to determine the target speed. In the absence of accepted criteria, three equal speed ranges are applied for the purposes of this section. Thus, the operation is defined as uncongested if the speed is above two-thirds of the target speed. It is defined as severely constrained when the speed is below one-third of the target speed, and it is considered congested in the middle speed range. This stratification is used to produce performance measures directly (e.g., percent of time severely constrained). It is also used in computing other performance measures (e.g., release from a queue).

Computational Procedures for Stop-Related Measures

The two main stop-related measures are number of stops and stopped delay. The beginning of a stop is defined in the same way for both measures. The end of a stop is treated differently for stopped delay and number of stops. For stopped delay, the end of a stop is established as soon as the vehicle starts to move (i.e., its speed reaches 5 mi/h or greater). For determining the number of stops, some hysteresis is required. For purposes of this section, after a vehicle is stopped a subsequent stop is not recognized until it leaves the severely constrained state (i.e., its speed reaches one-third of the target speed).

Because subsequent stops are generally made from a lower speed, they can be expected to have a smaller impact on driver perception, operating costs, and safety. Recognizing this fact, the NCHRP 3-85 project proposed a "proportional stop" concept (*1*), in which the proportion of a subsequent stop is based on the relative kinetic energy loss and is therefore proportional to the square of the speed from which the stop was made. Thus, each time a vehicle speed drops below 5 mi/h, the number of stops is incremented by $(S_{max}/S_{target})^2$, where S_{max} is the maximum speed attained since the last stop and S_{target} is the target speed.

While this procedure has not been applied in practice, it is mentioned here because it offers an interesting possibility for the use of simulation to produce measures that could be obtained in the field but could not be estimated by the macroscopic deterministic models described in the HCM. The procedure is illustrated by an example later in this section.

Computational Procedures for Delay-Related Measures

The procedures for computing delay from vehicle trajectories involve aggregating all delay measures over each time step. Therefore, the results take the form of aggregated delay and not unit delay, as defined in Chapter 7. To determine unit delays, the aggregated delays must be divided by the number of vehicles involved in the aggregation. Partial trips made over a segment during the time period add some complexity to the unit delay computations.

The following procedures should be used to compute the various delay-related measures from vehicle trajectories:

- *Time step delay*: The delay on any time step is, by definition, the length of the time step minus the time it would have taken the vehicle to cover the distance traveled in the step at the target speed. This value is easily determined and is the basis for the remainder of the delay computations.
- *Segment delay*: Segment delay is represented by the time actually taken to traverse a segment minus the time it would have taken to traverse the segment at the target speed. The segment delay on any step is equal to the time step delay. Segment delays accumulated over all time steps in which a vehicle is present on the segment represent the segment delay for that vehicle.
- *Queue delay*: Queue delay is equal to the time step delay on any step in which the vehicle is in a queued state; otherwise, it is zero. Queue delays are accumulated over all time steps while the vehicle is in a queue.
- *Stopped delay:* Stopped delay is equal to the time step delay on any step in which the vehicle is in a stopped state; otherwise, it is zero. Because a vehicle is considered to be "stopped" if it is traveling at less than a threshold speed, a consistent definition of stopped delay requires that the travel time at the target speed be subtracted. Time step delays accumulated over all time steps in which the vehicle was in the stopped state represent the stopped delay.
- *Control delay*: Control delay is the additional travel time caused by operation of a traffic control device. It cannot be computed directly from the vehicle trajectories in a manner consistent with the procedures given in Chapters 18 and 31 for signalized intersection analysis. It is, however, an important measure because it is the basis for determining the level of service on a signalized approach.

Queue delay computed from trajectory analysis provides the most appropriate representation of control delay.

The queue delay computed from vehicle trajectories provides a reasonable approximation of control delay when the following conditions are met:

1. The queue delay is caused by a traffic control device, and
2. The identification of the queued state is consistent with the definitions provided in this section.

Computational Procedures for Queue-Related Measures

Procedures for computing queue-related measures begin with determining whether each vehicle in a segment is in a queued state. A vehicle is in a queued state if it has entered a queue and has not yet left it. The beginning of a queued state occurs when

- The gap between a vehicle and its leader is less than or equal to 20 ft,
- The vehicle speed is greater than or equal to the leader speed, and
- The vehicle speed is less than or equal to one-third of the target speed (i.e., the speed is severely constrained).

A separate case must be created to accommodate the first vehicle to arrive at the stop line. If the link is controlled (interrupted-flow case), the beginning of the queued state also occurs when

- No leader is present on the link,
- The vehicle is within 50 ft of the stop line, and
- The vehicle is decelerating or has stopped.

These rules have been found to cover all the conditions encountered.

The ending of the queued state also requires some rules. For most purposes, the vehicle should be considered to remain in the queue until it leaves the link. The analysis is done on a link-by-link basis. In the case of queues that extend over multiple links, a vehicle leaving a link immediately enters the queue on the next link. Experience with trajectory analysis has shown that other conditions need to be applied to supplement this rule. Thus, the end of the queued state also occurs when

- The vehicle has reached two-thirds of the target speed (i.e., uncongested operation), and
- The leader speed is greater than or equal to the vehicle speed or the vehicle has no leader in the same link.

The additional conditions cover situations in which, for example, a vehicle escapes a queue by changing lanes into an uncongested lane (e.g., through vehicle caught temporarily in a turn bay overflow).

Chapters 18 and 31 offer the following guidance on estimating queue length:

1. The maximum queue reach (i.e., back of queue, or BOQ) is a more useful measure than the number of vehicles in the queue, because the BOQ causes blockage of lanes. The maximum BOQ is reached when the queue has almost dissipated (i.e., has zero vehicles remaining).

2. A procedure is prescribed to estimate average maximum BOQ on a signalized approach.

The BOQ at any time step will be determined by the position of the last queued vehicle on the link plus the length of that vehicle.

Because of its macroscopic nature, the HCM queue estimation procedure cannot be applied directly to simulation. On the other hand, simulation can produce additional useful measures because of its higher level of detail. The first step in queue length determination has already been dealt with by setting up the rules for determining the conditions that indicate when a vehicle is in a queue. The next step is to determine the position of the last vehicle in the queue.

The BOQ on any step is a relatively simple thing to determine. The trick is to figure out how to accumulate the individual BOQ measures over the entire period. Several measures can be produced.

1. The maximum BOQ at some percentile value—for example, 95%;
2. The maximum BOQ on any cycle at some percentile value—for example, 95%;
3. The historical maximum BOQ (i.e., the longest queue recorded during the period;
4. The probability that a queue will back up beyond a specified point; and
5. The proportion of time that the queue will be backed up beyond a specified point.

Some of these measures are illustrated later in an example.

Computational Procedures for Density-Related Measures

The uninterrupted-flow procedures described in the HCM base their level of service estimates on the density of traffic in terms of passenger cars per mile per lane (pc/mi/ln). In one case (freeway merges and diverges), the density is estimated only for the two lanes adjacent to the ramp.

Density computations do not require a detailed analysis of the trajectory of each vehicle. They are best made by simply counting the number of vehicles in each lane on a given segment, recognizing that the results represent actual vehicles and not passenger cars.

For comparable results, it essential to convert the simulated densities to pc/mi/ln, especially if simulation tools are used to evaluate the level of service on a segment. Because the effect of heavy vehicles on the flow of traffic is treated microscopically, there is no notion of passenger car equivalence in simulation modeling. In addition, traffic flow models may differ among the various simulation tools in their detailed treatment of heavy vehicles. Therefore, it is not possible to prescribe a simple conversion process that will ensure full compatibility with the HCM's level-of-service estimation procedures. One possible method to develop passenger-car-equivalence conversion factors involves multiple simulation runs:

1. Use the known demand flow rates, v, and truck proportions to obtain the resulting segment density in vehicles per mile per lane (veh/mi/ln), d_1.
2. Use the known demand flow rates, v, with passenger cars only to obtain the resulting segment density in veh/mi/ln, d_2.

3. Determine the heavy-vehicle equivalence factor as $f_{HV} = d_2/d_1$.
4. Set the demand flow rates to v/f_{HV} with passenger cars only to obtain the resulting segment density in pc/mi/ln.

This process is more precise because it adheres to the definition of passenger car equivalence. Unfortunately, it is too complicated to be of much practical value. There are, however, two methods that could produce a more practical approximation. Both require determining the heavy-vehicle adjustment factor, f_{HV}, by the method prescribed in Chapter 11 for basic freeway segments. This method is also referenced and used in the procedural chapters covering other types of freeway segments. The simplest approximation may be obtained by running the simulation with known demand flow rates and truck proportions and then dividing the simulated density by f_{HV}. Another approximation involves dividing the demand flow rates by f_{HV} before running the simulation with passenger cars only. The resulting densities are then expressed in pc/mi/ln. The second method conforms better to the procedures prescribed in Chapters 11 to 13, but the first method is probably easier to apply.

Follower density is an emerging density-based measure for two-lane highways (*2*, *3*). It is defined as the number of followers per mile per lane. A vehicle can be classified as following when

- The gap between the rear and the front ends of the leading and following vehicles, respectively, are shorter than or equal to 3 s, and
- The speed of the following vehicle is not more than 12 mi/h lower than that of the preceding vehicle.

The follower density can be derived from point measurements by means of the following formula:

$$\text{Follower Density} = \text{\% Followers} \times \text{Flow Rate}/\text{Time Mean Speed}$$

Although this performance measure is not computed by the procedures in the HCM, it is worth mentioning because it has attracted significant international interest and is a measure that can easily be computed by vehicle trajectory analysis.

Analysis of a Signalized Approach

The simple approach to a signalized intersection (Exhibit 24-2) is now converted to a two-lane approach with a length of 2,000 ft. A 10-min (600-s) analysis period is used. The cycle length is 60 s, giving 10 cycles for inspection. The analysis period would normally be longer, but 10 min is adequate for demonstration purposes.

Trajectory Plots

The trajectory plot for the first few cycles is shown in Exhibit 24-20. The vehicle track selected for later analysis is also shown in this exhibit.

Exhibit 24-20
Trajectories for Several Cycles on a Signalized Approach

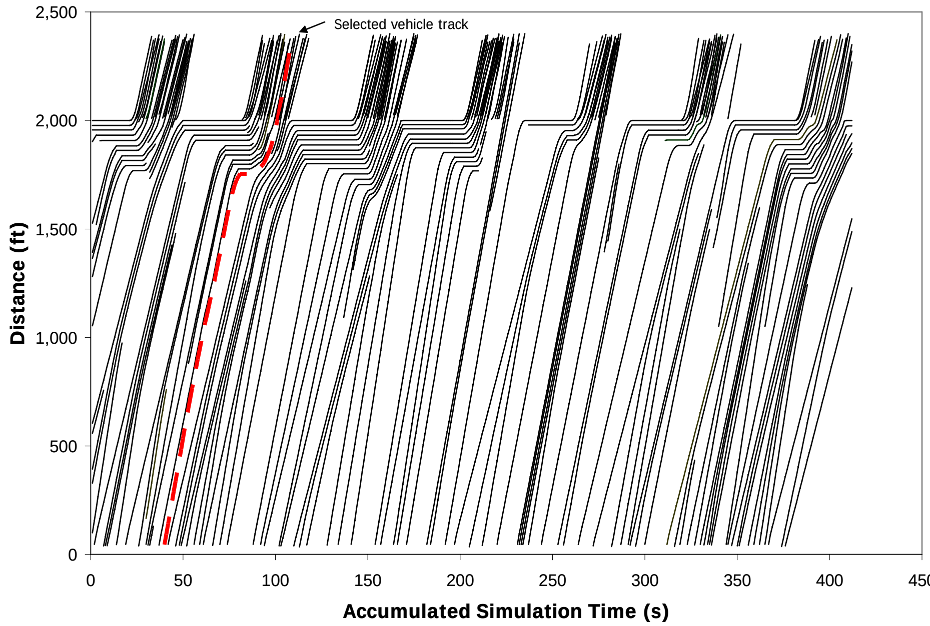

Two individual trajectory analysis plots are shown in Exhibit 24-21. The first plot shows the trajectories of two vehicles where the progress of the subject vehicle is constrained by its leader. The second plot shows the speed and acceleration profiles for the subject vehicle.

Exhibit 24-21
Example Trajectory Analysis Plots

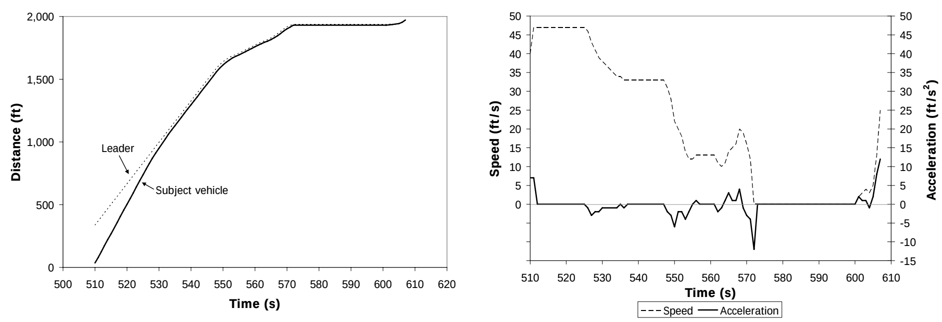

(a) Subject Vehicle and Leader Vehicle Trajectories

(b) Speed and Acceleration Profile of Subject Vehicle

Analysis of Stops

An example of the analysis of a single vehicle selected from the entire trajectory plot is shown in Exhibit 24-22. With the definition of a partial stop based on the NCHRP 3-85 kinetic energy loss concept, the total stop value was 1.81 because the second stop was made from a lower speed.

Exhibit 24-22
Analysis of a Full and Partial Stop

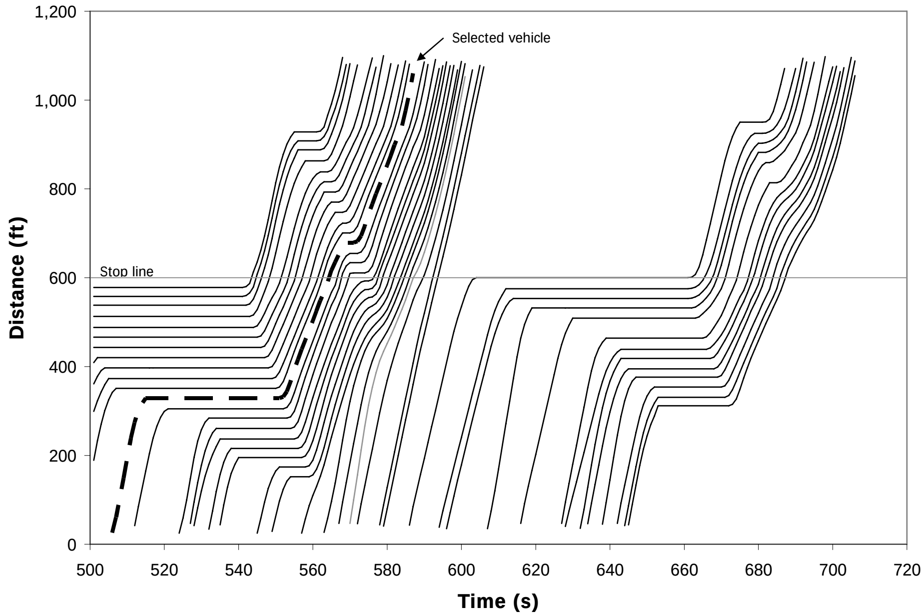

(a) Vehicle Trajectories

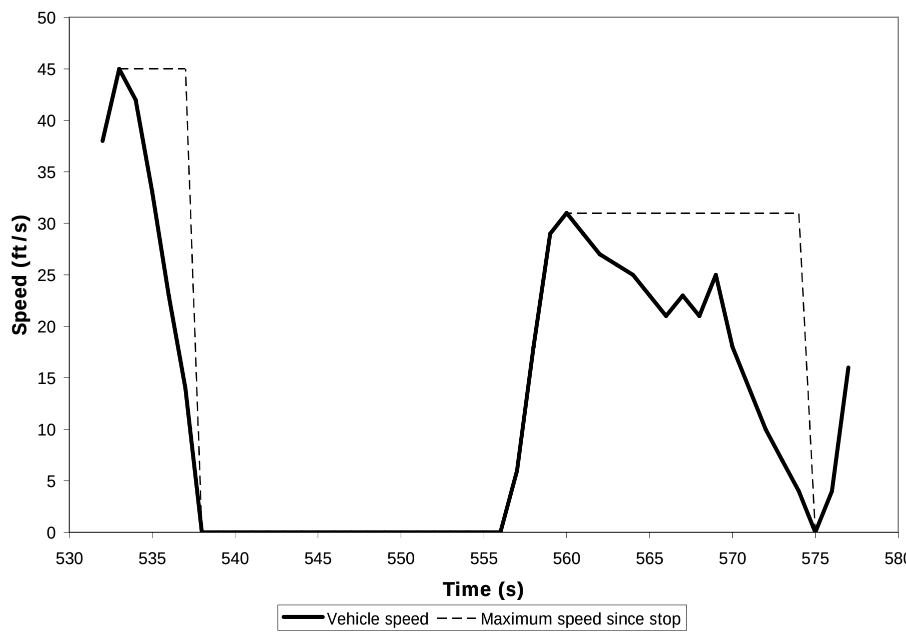

(b) Selected Vehicle Speed

Segment delay	Queue delay	Stop delay	Number of stops
34.64 s	33.23 s	20 s	1.81

(c) Performance Measures for Selected Vehicle

Queuing Analysis

Exhibit 24-23(a) illustrates the queue length (BOQ) per step for one lane of the signalized approach over all the time steps in the period. The 10 cycles are discernible in this figure. Also, a considerable variation in the cyclical maximum BOQ is evident.

It is important to distinguish between the percentile instantaneous BOQ and the percentile maximum BOQ per cycle. For the instantaneous BOQ, the individual observation is the BOQ on any step. So the sample size is the number of steps covered (600 in this case). For cyclical maximum BOQ, the individual observation is the maximum BOQ in any cycle, so the sample size is the number of cycles (10 in this case). The maximum BOQ in any cycle can be determined only by inspecting the plotted instantaneous values. No procedure is proposed here for automatic extraction of the maximum cyclical BOQ from the instantaneous BOQ data.

A statistical analysis showing the average BOQ, the 95th percentile BOQ (based on 2 standard deviations past the average value), and the historical maximum BOQ is presented in Exhibit 24-23(b). One important question is whether the 95% BOQ can be represented statistically on the basis of the standard deviation, assuming a normal distribution. The BOQ histogram showing the distribution of instantaneous BOQ for the 600 observations is shown in Exhibit 24-24. The appearance of this histogram does not suggest any analytical distribution; however, the relationship between the 95% BOQ and the historical maximum appears to be reasonable for this example.

Exhibit 24-23
Back of Queue Analysis by Time Step

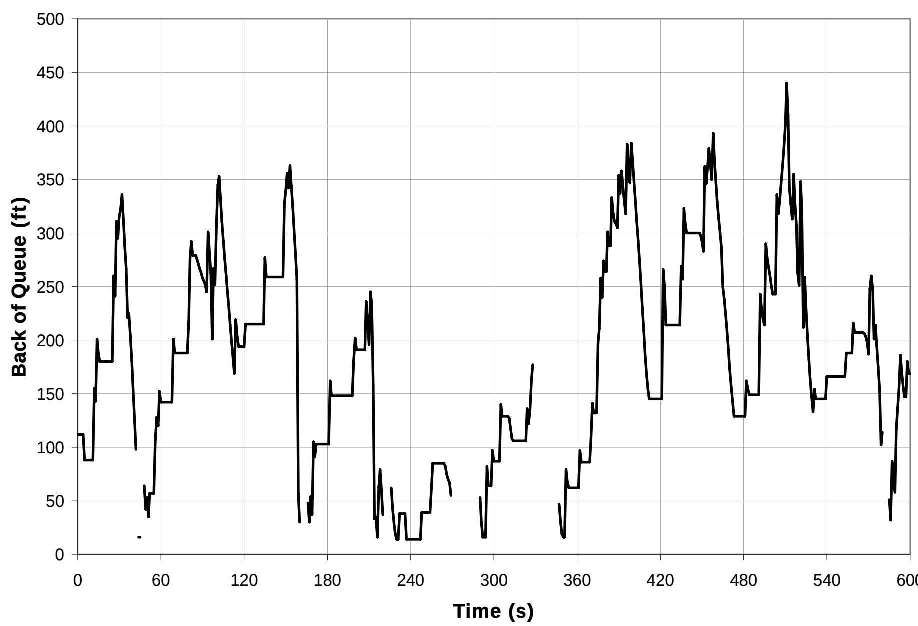

(a) Back of Queue Plot

Average queue	Standard deviation	95th percentile queue	Maximum queue
174 ft	110 ft	395 ft	440 ft

(b) Queue-Related Performance Measures

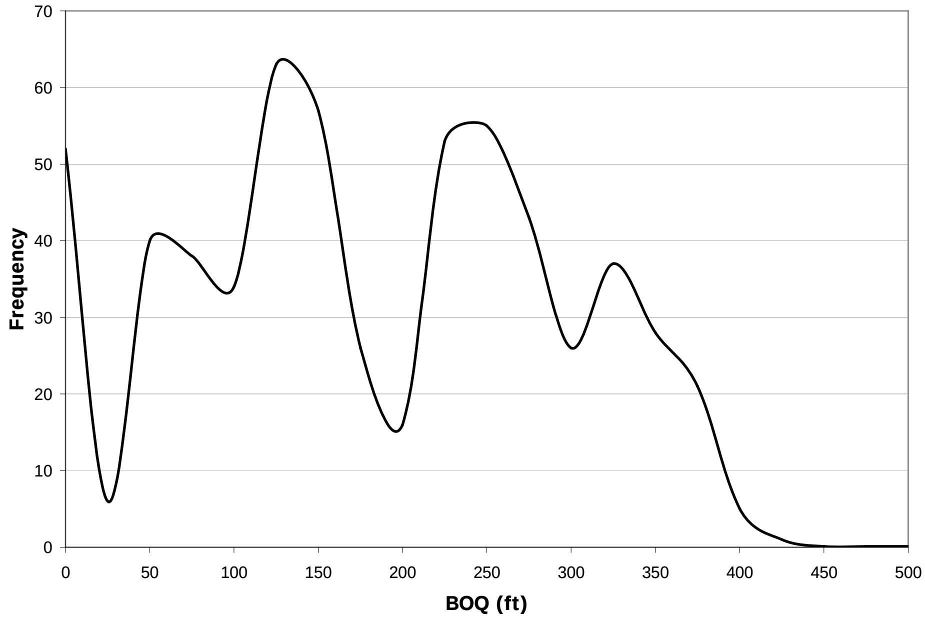

Exhibit 24-24
Back of Queue Histogram

The queue length on an isolated approach that is approaching saturation will have a near uniform distribution (i.e., equal probability of all lengths between zero and the maximum). The standard deviation of a uniform distribution is greater than one-half of the mean, so the 95th percentile estimator (mean value plus 2 standard deviations) will be greater than the maximum value. This situation raises some doubt about the validity of basing the 95th percentile BOQ on the standard deviation, especially with cyclical queuing.

Delay Analysis for a Single Trajectory

A comparison of the accumulated delay by all definitions for the selected vehicle track indicated in Exhibit 24-20 is presented in Exhibit 24-25(a). The relationships between segment delay, queue delay, and stopped delay are evident in this figure. The segment delay begins to accumulate before the vehicle approaches the intersection because of mid-segment interactions that reduce the speed below the target speed. The queue delay begins to accumulate as the vehicle enters the queue, and the stopped delay begins to accumulate a few seconds later. The stopped delay ceases to accumulate as soon as the vehicle starts to move, but the queue delay continues to accumulate until the vehicle leaves the link.

The time step delay analysis plots shown in Exhibit 24-25(b) provide additional insight into the operation. The time step delay is close to zero as the vehicle enters the segment, indicating that the speed is close to the target speed. Small delays begin to accumulate in advance of the intersection. The accumulation becomes more rapid when the vehicle enters the queue. The periods when the vehicle is in the stopped and the queued state are also shown in this figure.

Exhibit 24-25
Accumulated Delay by Various Definitions

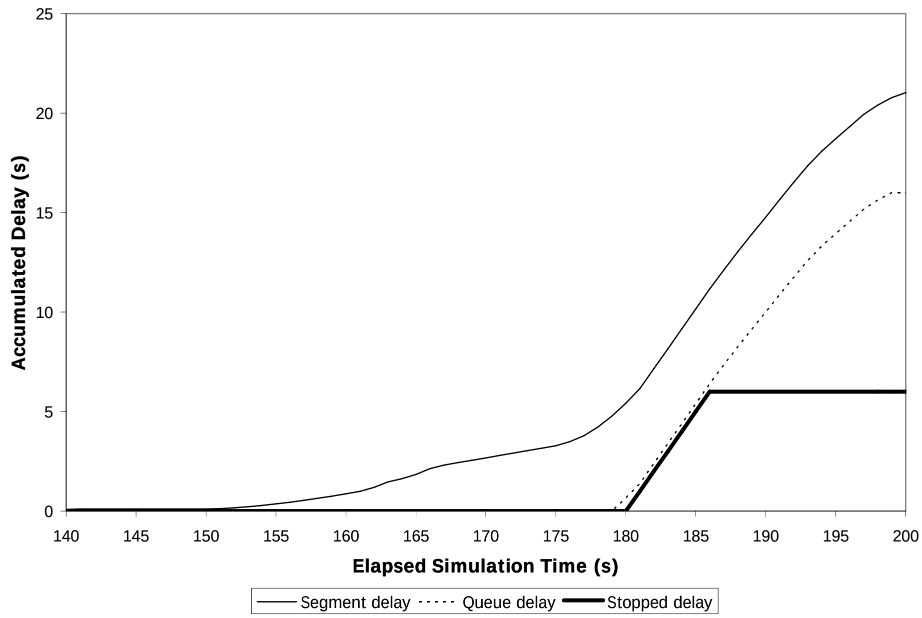

(a) Accumulated Delay

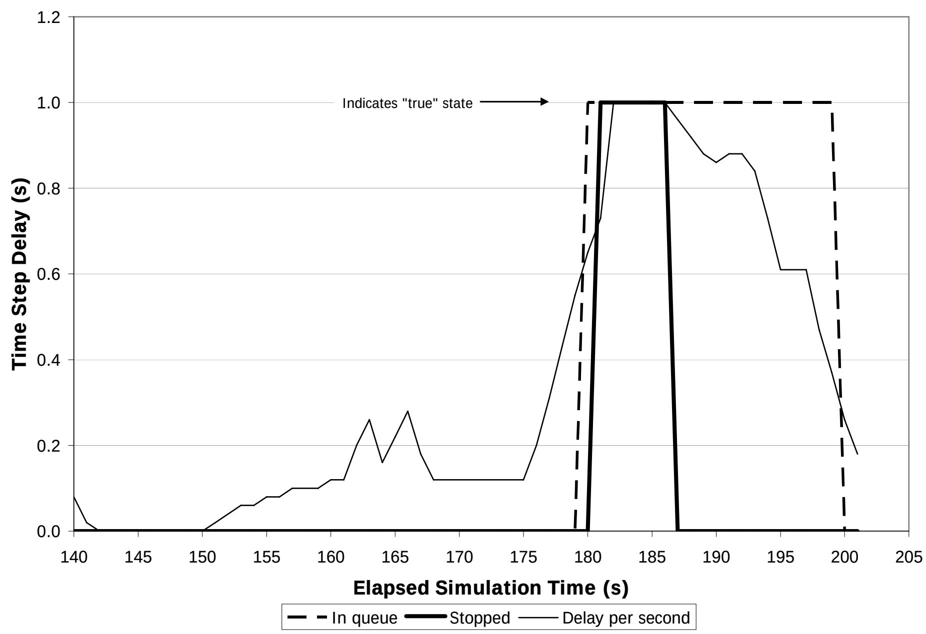

(b) Time Step Delay

It was indicated previously that the value of control delay cannot be determined by simulation in a manner that is comparable with the procedures prescribed in Chapters 18 and 31. Because this segment terminates at a signal, it is suggested that the queue delay would provide a reasonable estimate of control delay because the queue delay offers a close approximation to the delay that would be measured in the field.

Delay Analysis for All Vehicles on the Segment

The preceding example dealt with accumulated delay of a single vehicle traversing the segment. A useful delay measure requires the accumulation of delay to all vehicles traversing the segment during the period. An example is shown in Exhibit 24-26. In keeping with the recommendations offered elsewhere (*4*), only vehicles that traversed the entire link during the period are included in this analysis. Therefore, the number of vehicles analyzed (210) is lower than the number of vehicles that were actually on the link during the period (286).

Exhibit 24-26
Delay Analysis for All Vehicles on a Segment

	Segment Delay (s)	Queue Delay (s)	Stop Delay (s)	No. of Stops
Lane 1	3,128	2,562	1,957	95.4
Lane 2	3,400	2,793	2,047	96.2
Total	6,529	5,355	4,004	191.6
Average per vehicle	31.09	25.50	19.07	0.91

Analysis of a Freeway Segment

A performance analysis of the freeway merge area originally shown in Exhibit 24-9 is presented here. A single vehicle is selected from the trajectory plot and its trajectory is analyzed. The results are shown in Exhibit 24-27. The analysis produced segment delay and queue delay. This segment was very congested, as indicated by the trajectory plot. No stopped delay was produced because the vehicle never actually came to a stop (i.e., its speed stayed above 5 mi/h).

Exhibit 24-27
Longitudinal Analysis of Delay for a Selected Vehicle in a Merge Area

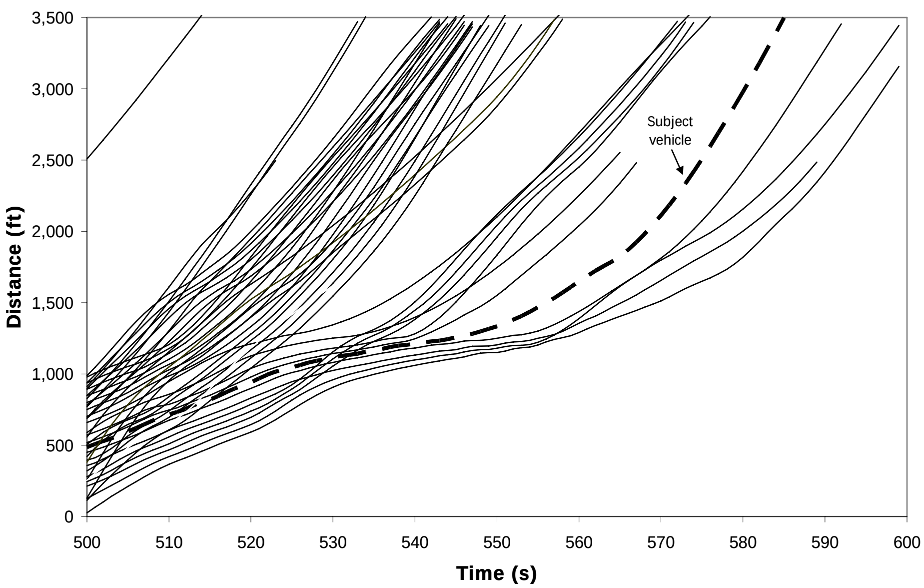

(a) Vehicle Trajectories

Segment delay	Queue delay	Stopped delay
39.58 s	37.01 s	0 s

(b) Delay-Related Performance Measures for Subject Vehicle

A spatial analysis of the entire segment can also be performed to produce the following measures by lane:

- Average density over the segment,
- Percent slow vehicles (i.e., traveling at less than two-thirds the target speed),
- Percent queued vehicles,
- Average queue length (measured from front of queue to BOQ),
- Average BOQ position,
- Maximum BOQ position, and
- Percent of time steps when the queue overflowed the segment.

The results are presented in tabular form in Exhibit 24-28. The values are presented by lane and are combined for Lanes 1 and 2 for compatibility with the HCM definition of merge area density.

Exhibit 24-28
Example Spatial Analysis by Lane

	Lane 1	Lane 2	Lane 3	Acceleration Lane
Average density (veh/mi/ln)	73.4	51.0	43.6	9.9
Percent slow vehicles (%)	88.4	68.5	41.5	65.7
Percent queued vehicles (%)	63.4	22.0	2.4	26.7
Average queue length (ft)	600	215	15	40
Average back of queue (ft)	1,471	1,119	135	562
Maximum back of queue (ft)	1,497	1,497	1,492	1,474
Percent overflow	66.1	29.6	0.5	0.17

Note: Average Lane 1 and Lane 2 density is 62.2 veh/mi/ln.

2. REFERENCES

Many of these references can be found in the Technical Reference Library in Volume 4.

1. Courage, K. G., S. Washburn, L. Elefteriadou, and D. Nam. *Guidance for the Use of Alternative Traffic Analysis Tools in Highway Capacity Analyses.* National Cooperative Highway Research Program Project 3-85 Final Report. University of Florida, Gainesville, 2010.
2. Van As, S. C., and A. Van Niekerk. The Operational Analysis of Two-Lane Rural Highways. Presented at 23rd Annual Southern African Transport Conference, Pretoria, South Africa, July 2004.
3. Catbagan, J. L., and H. Nakamura. Probability-Based Follower Identification in Two-Lane Highways. Presented at 88th Annual Meeting of the Transportation Research Board, Washington, D.C., 2009.
4. Dowling, R. *Traffic Analysis Toolbox Volume VI: Definition, Interpretation, and Calculation of Traffic Analysis Tools Measures of Effectiveness*. Federal Highway Administration, Washington, D.C., 2007.