4. Traffic Network Analysis with NETSIM--A User Guide. Implementation Package FHWA-IP-80-3. FHWA, U.S. Department of Transportation, 1980.
5. F.B. Lin. Optimal Timing Settings and Detector Lengths of Presence-Mode Full-Actuated Control. Draft Final Report. Office of University Research, U.S. Department of Transportation, Dec. 1984.
6. Development and Testing of INTRAS, A Microscopic Freeway Simulation Model, Vol. 1. Report FHWA/RD-80/106. FHWA, U.S. Department of Transportation, 1980.
7. Network Flow Simulation for Urban Traffic Control System--Phase II, Vol. 5. FHWA-RD-77-45. FHWA, U.S. Department of Transportation, 1977.

Publication of this paper sponsored by Committee on Traffic Flow Theory and Characteristics.

Another Look at Identifying Speed-Flow Relationships on Freeways

BRIAN L. ALLEN, FRED L. HALL, and MARGOT A. GUNTER

ABSTRACT

Despite approximately 50 years of research on highway operating characteristics, the way in which the speed-flow relationship moves between free flow and congested flow conditions is still not clearly understood. The speed-flow relationship as it pertains to those transitions is investigated using an extensive data set collected on the Queen Elizabeth Way freeway in Ontario. Two different analytical approaches are used: time-connected plots of mean speed and mean flow and an event-based trace, averaged with respect to the transition to and from congested flow. The results confirm several aspects of conventional understanding but also raise questions that are hard to answer with the conventional interpretations of speed-flow relationship.

Highway capacity and the operational characteristics of uninterrupted traffic flow have been the focus of research for at least the past 50 years. Since Greenshields' work in the 1930s, traffic researchers have devoted considerable attention to investigating and interpreting the fundamental characteristics of traffic flow on freeways. The dozens of research papers produced during those 50 years have typically documented either the results of interpretive empirical studies or the degree of success achieved in relating those results to known or proposed theoretical concepts (models). Certainly it is a well-known and well-researched area. Why then another paper on this subject?

The answer consists of three complementary parts. First, the subject area itself remains relevant. As freeway systems become more complex and experience ever-increasing traffic loads, the effective management of those systems becomes increasingly important. To manage them effectively, operating agencies must have reliable information about traffic flow on which to base appropriate actions. This is true whether overall system management or more specific and data-demanding activities, such as entrance ramp control or incident detection and response, are considered.

Second, it is well accepted that the representation of speed-flow relationships first presented in the Highway Capacity Manual (HCM) (1), and more recently updated in the Interim Materials on Highway Capacity (2), cannot completely reflect all actual conditions witnessed on the many different existing freeway systems. There is little doubt that at the very least a further updating of those curves is required to reflect current operating characteristics, even if the only result is simply to "calibrate," on a site-specific basis, the traditional approach to highway capacity and level of service.

This leads to the third part of the answer. In trying to calibrate the traditional approach with their data, many traffic engineers have had difficulty with the traditional approach to describing the speed-flow relationship as a smooth, continuous curve as depicted in the HCM material and the vast majority of standard traffic engineering references. Some researchers, suspecting that the curves are not in fact continuous and perhaps not always smooth, have proposed other solutions. Perhaps it is appropriate even after 50 years of research to take yet another look at precisely what does happen with traffic flow on freeways.

In this paper only the relationship between speed and flow on freeways, particularly as the flow moves between the congested and uncongested regimes, is discussed. The next section provides some background for the analysis by including some of the expressions of dissatisfaction with current understanding of traffic flow characteristics. That discussion also helps to focus on possible analytical techniques by suggesting why other analysts are not satisfied with their results. In the subsequent sections the data set used is described, the results of the analysis are presented, and the conclusions that appear to follow from the analysis are provided.

BACKGROUND

Most researchers would surely agree that there are three primary issues related to understanding freeway traffic flow relationships:

1. Identification of the basic variables.
2. Formulation of the fundamental interactions among these variables. This includes consideration of functional relationships that describe the nature of uninterrupted flow throughout the range of low volume-to-capacity ratios (V/C), high V/C, congested operation, and transitions between them.
3. Quantification of the flow relationships, that is, identification of the numerical magnitudes of ideal lane capacity, free flow speed, and so on.

There is certainly agreement regarding the first issue: speed, flow, and density are the variables of interest. Some concern has been expressed in the past over the use of particular speed measures, but that debate appears to have been practically resolved with the more recent adoption of average (running) speed. This will be the speed referred to in the remainder of this paper. Similarly, the vast majority of researchers use short duration volume counts expanded to hourly flow rates. Again, there has been debate over how short the time durations should be, but generally it is accepted that the shorter durations (1, 2, or 5 min) are best suited for investigations of flow interactions and characteristics. Because the present work is concerned exclusively with identifying such interactions, 5-min volumes and resultant hourly flow rates will be used throughout. Finally, the usefulness of density or occupancy rates cannot be denied. The measure is appropriate to level of service concerns and helpful in identifying basic relationships. Because this work is concerned primarily with speed-flow relationships, however, few comments will be made regarding density.

Although the second issue of formulating fundamental relationships has not been as readily resolved, there are important areas of agreement. Most traffic engineers would agree that the basic shape is somewhat similar in nature to the curves shown in Figure 1. In particular, it is accepted that traffic operates in either an uncongested state on the upper branch of the curve or in a congested state on the lower branch of the curve. This basic form was well represented by the HCM curve (Curve 1) in the figure (1). There is also recent agreement that speeds on the upper branch remain relatively constant over much larger ranges of flow rate than previously depicted, as typified by the other plots (Curves 2-4) in Figure 1 (2-4). In addition, there is agreement that congested operations occur approximately as shown by the lower branch in the figure. There is less agreement on the way in which the upper and lower branches are joined and on the operational nature of transitions between the branches. It is possible that these transitions are dependent

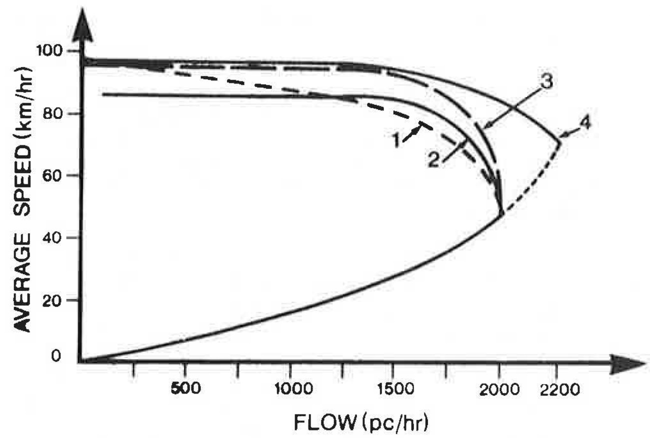

FIGURE 1 Speed-flow relationships from various studies: 1, 1965 Highway Capacity Manual *(1)*; 2, Interim Materials on Highway Capacity *(2)*; 3, Ontario 1978 *(3)*; and 4, Ontario 1983 *(4)*.

on demand and capacity characteristics and location with respect to bottleneck sections. The third issue of quantification is a relatively simple statistical one when the form of the relationship has been resolved. Unfortunately, the shape of the relationship at high flow rates has not been satisfactorily resolved, so magnitude estimation is still a problem.

Most practitioners believed that concerns expressed following publication of the 1950 HCM (5) had been satisfactorily addressed with the revised notions presented in the 1965 version. The speed-flow relationships presented in the 1965 HCM remained at the forefront of acceptance and use for many years. However, as freeway design, operations, and control expertise grew during the late 1960s and the 1970s, researchers and practitioners recognized that, for whatever reasons, the HCM representations were outdated. Despite this recognition, the results of a major 1977 literature review on flow relationships indicated that the published studies "tend to be quite detailed, but for limited sections of highway, and none attempts to generalize the basic relationships" (6,p.3). In other words, although there was dissatisfaction and a need to revise the old understanding, no one had yet provided a new interpretation. It is particularly interesting to note in that review that Roess could find only 15 references that would "probably be useful in the development of narrative and background portions" (6,p.2), and none of these were considered to be of great utility in revising existing relationships. Of the 15, one was a Greenshields' paper published in the 1930s, and five others had been published before the 1965 HCM.

Other researchers have continued to recognize the problem in work published in the 1980s, as exemplified by such comments as the following from Koshi et al. (7,p.403).

> This paper deals with vehicular flow characteristics especially in a congested region, and attempts to describe what really happens and why.
>
> It has been pointed out that vehicular flows oscillate in congested conditions and that there is a discrepancy in the speed-density-volume relationships between free flow and congested flow regions. The phenomena, however, seem not to have been explored thoroughly enough to understand the total picture of vehicular flow characteristics.

Even more recently, Hurdle again voiced the recurrent concern (8,p.127):

> A good understanding of the way in which speed varies with flow is an essential prerequisite to the creation and use of any level-of-service concept for freeways. Unfortunately, misinformation about this relationship abounds.

Presumably the dissatisfaction over the current level of understanding refers to the speed-flow interactions when flow rates are extremely high, or when transitions are made from upper-branch to lower-branch operation and back again, or both. It would appear that the basic nature of the problem facing those researchers lies either with the characteristics of the data that are available to them or with the analysis and interpretation procedures used, or some combination of the two. Although there have been some problems in the past, current technology permits collection of adequate data sets with little difficulty. As a result, the problem appears to lie with the choice of analytical procedures. This is normally dependent on the paradigm selected or hypotheses to be tested. For example, if it were assumed that speed and flow vary according to an HCM-type curve, data could be plotted and standard curve-fitting techniques adopted to generate the entire two-branch relationship. Many researchers have attempted to refine this approach, at least using speed and density data, by fitting curves to the uncongested and congested regimes separately (9-11), achieving some degree of explanatory success. Because there is no doubt about the existence of two regimes of traffic flow, that surely should also have a bearing on the manner in which speed-flow data are analyzed.

To illustrate what such an approach means for analysis, consider the attempts of Mahabir (12) in 1980 to fit a relationship to the rather extensive data set that is used in this study as well. Given a typical data set as shown in Figure 2, he assumed a two-regime model and used a number of regression techniques to try to determine the curve.

The distinction between the congested and uncongested states was made by arbitrarily selecting a speed at which he felt capacity operation occurred. It appears to be just as reasonable to fit the curve by eye, as in fact Mahabir ended up doing, because either way prejudgment of parameters, even on the already-assumed curve, must be made. Hurdle and Datta (8) also discuss the problem of curve fitting at or approaching capacity, and present five curves, all compatible with their data, that could represent the speed-flow relationship in this area. Although three of the curves would be more generally accepted, two highly unusual curves fit the data just as well.

In light of all the uncertainty, how much reliance can be placed on the traditional understanding? Traffic engineers during the past 50 years have repeatedly stated that the speed-flow relationship is not fully understood. Therefore, it may be worthwhile to take a different analytic approach to the problem. The procedures adopted for this investigation represent some different approaches and are discussed following a brief description of the data set.

SPEED-FLOW DATA

As with any analytic exercise directed toward determining the "true" nature of a real-world phenomena from only a sample of data, the researcher must have confidence in that data set, or at least know its limitations. For many years, researchers were hampered by relatively small data samples with consequent low levels of confidence as to their representativeness. During the past 20 years, however, the increasing implementation of freeway surveillance and control systems and the availability of high technology portable equipment have allowed collection of vast amounts of data with relatively high reliability. The problems related to confidence have subsequently been largely removed, leaving only the question of how best to organize and analyze these large data samples.

The data used in this study, originally obtained and analyzed by Mahabir in 1980, come from the Ontario Ministry of Transportation and Communications Freeway Surveillance and Control System (13). The data were collected in 1979-1980, and at that time the system operated on a 5-km section of the Queen Elizabeth Way (QEW) between Oakville and Toronto, where morning commuter traffic created congested flow conditions on the three eastbound lanes (Figure 3). The system comprised nine mainline detector stations with induction loop pairs in each of the three directional lanes, ramp metering on five entrance ramps, and closed-circuit television (CCTV) surveillance cameras operated from a control center. The limiting capacity restriction (bottleneck) was downstream of Station 9 where heavy Highway 10 entrance ramp traffic merged with the three through lanes.

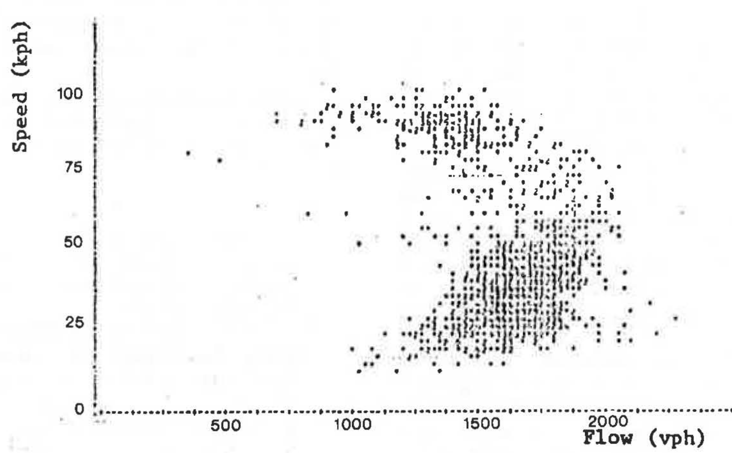

FIGURE 2 Scatterplot of data used by Mahabir (*12*, Figure A2.23, Station 9, middle lane).

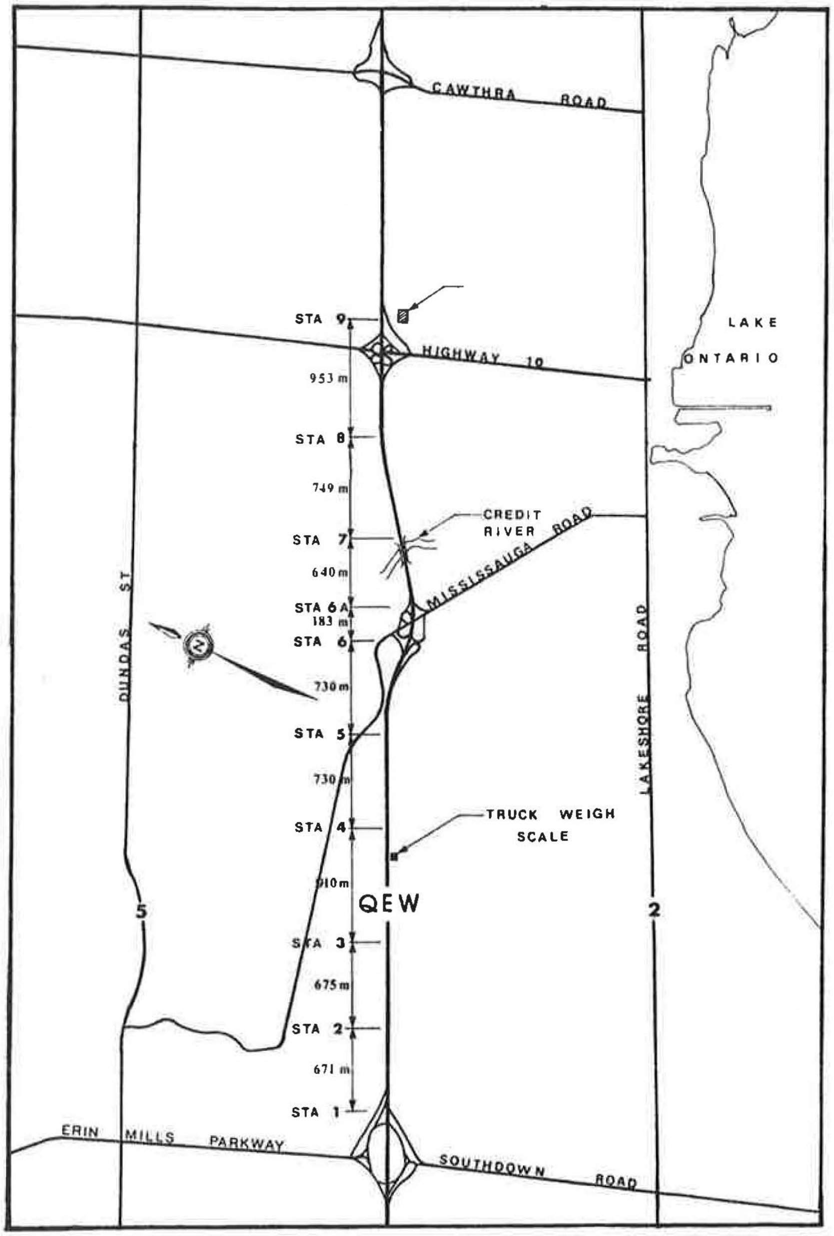

FIGURE 3 Station locations for the Freeway Surveillance and Control System, Queen Elizabeth Way, Ontario [reproduced from Case and Williams *(13)* Figure 3.4].

The data obtained from that system span 8 months from July 1979 to February 1980 and consist of approximately 2.5 hr of collection during weekday morning peak periods. From each pair of induction loops the following information was compiled for each lane for 5-min intervals: occupancy at the downstream loop, occupancy at the upstream loop, volume of vehicles longer than 7.6 m, total volume of vehicles, and average speed. The data were stored on magnetic tapes by the Data General computer that operated the system. In addition, a complete log of daily weather conditions and incidents (accidents, breakdowns, and so forth) was available.

Because the purpose of the current analysis is to identify the basic speed-flow characteristics, it was decided that only data representing "ideal" conditions would be used. Mahabir had already identified all days for which no incidents, accidents, or adverse weather had been logged; 68 such days were extracted from the 8-month period. This compilation conservatively represented only the "best" days of operation and has been used for all analyses reported in this paper.

The data were maintained and analyzed on a lane-by-lane basis because Mahabir's work suggested considerable differences in results according to lane. This distinction required that the analysis be performed in passenger car equivalents in order that the shoulder lane with 12 to 25 percent trucks, middle lane with 2 to 10 percent, and median lane with 0 to 2 percent could be easily compared. A passenger car equivalency of 2 was used, as recommended in TRB Circular 212 (2).

ANALYSIS

The appropriate analytical procedures to use with data such as these are not obvious. Mahabir's scatterplots (Figure 2) provide a useful picture of the range of likely and possible values of speed-flow points, but they do not provide an understanding of what actually occurs on the roadway. The present analysis, then, began with his results within the context of current speed-flow ideas and attempted to sort out the data to best describe freeway speed-flow behavior. For the most part, graphic techniques instead of statistical tests have been used. The reason for this is that statistical curve fitting is appropriate only when there is some theory to suggest the type of curve to use. The present question is simpler: what is the general pattern underlying the data?

Two types of analysis are reported here. The two procedures are presented in the order they were tried because the first analysis provided some insights and further questions that led to the second method of analysis. Other approaches were also attempted but do not merit discussion because the results added relatively little to understanding the pattern. (Those approaches were plotting speed versus time of day, and flow rates versus time of day for several days; and plotting the 68-day average speed and flow versus time.) In this discussion the focus will be mainly on Station 9 because it is situated closest to the capacity restriction and, therefore, experiences the longest period of congestion. However, there are few differences in the general trends between stations, and any of significance will be noted.

Time-Connected Plots of Mean Speeds Versus Mean Flows

The first analysis considered how speeds and flows are related over time. Such an approach has been used before to assist in resolving the problems of

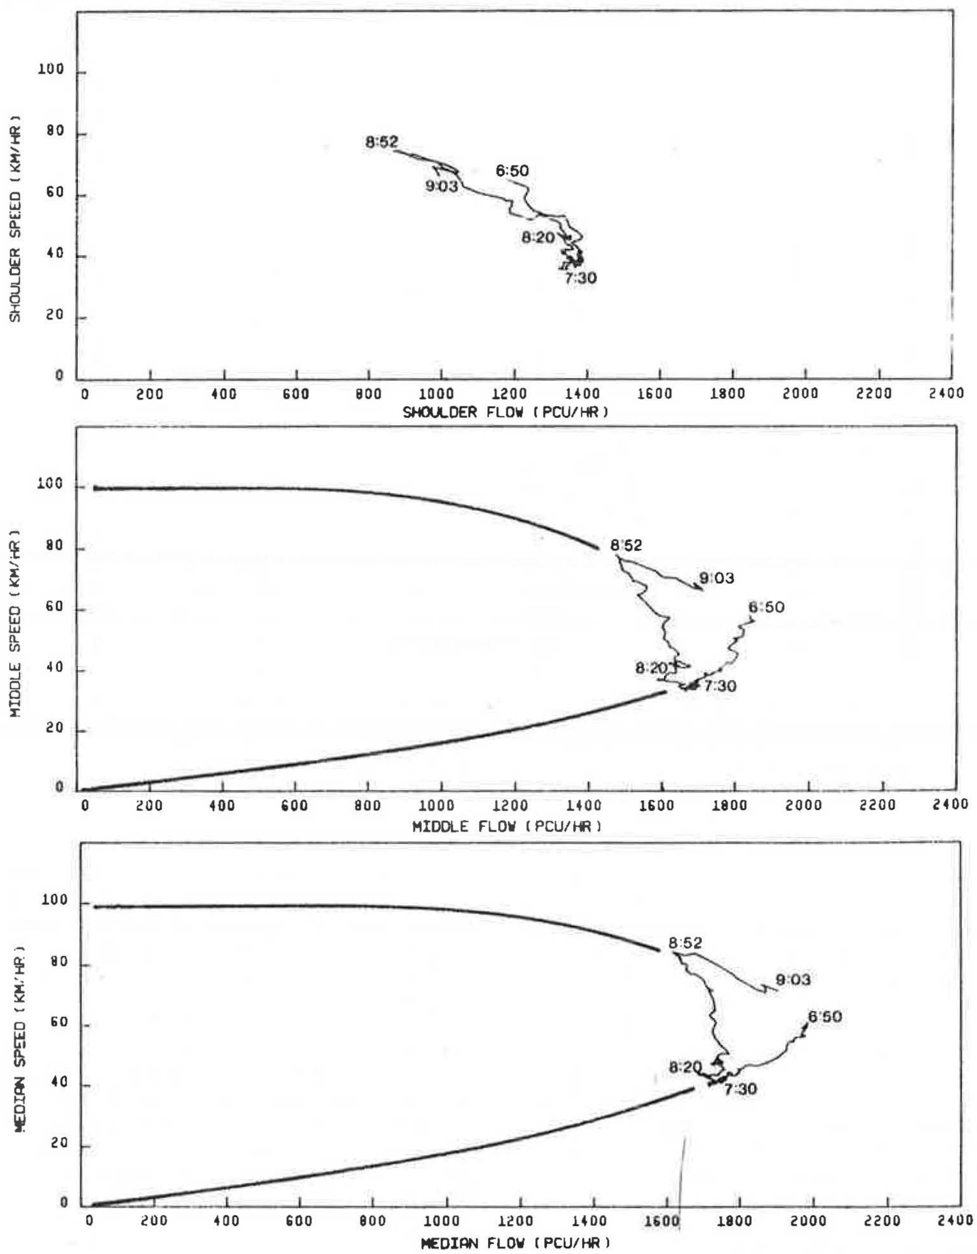

FIGURE 4 Station 9 mean speeds versus mean flow rates (over 68 days) connected by clock time; heavy lines suggest an HCM-type curve consistent with the data.

distinguishing between congested and uncongested speed-flow pairs, but only with daily data not averages over many days (14). Mean speeds and flows are plotted in the time-connected diagram of Figure 4. There are several interesting points to examine in the median and middle lane plots, which are quite similar in nature, particularly in relation to conventional interpretation. First, these data can be interpreted to be consistent with the standard HCM-type speed-flow representation as suggested by the heavy solid lines that have been positioned to fit these data. Second, the speed decrease takes a long time to occur (6:50 a.m. to 7:30 a.m.). If that decrease takes place along the lower branch of the suggested relationship, the duration is not too surprising. However, if operation at 6:50 a.m. was in the uncongested state, a more sudden drop in speeds at approximately the same flow rate would be expected. Because the speeds start at about 60 km/hr, acceptance of the lower-branch operation argument would seem reasonable. Third, regardless of the manner in which speeds and flows decreased, the relatively steady operation between 7:30 a.m. and 8:20 a.m. is to be expected. Queues and the resultant storage requirements imposed by the downstream bottleneck remain relatively constant during the congested period and appear to fluctuate less than 10 percent in speed and flow. Fourth, the speed increase between 8:20 a.m. and 8:52 a.m. appears reasonable in relation to the sketched curves, but detailed examination raises some questions.

The first question relates to duration. It is clear that the increase can be related to a transition from congested operation on the lower branch to uncongested operation on the upper branch. If so, why does the recovery take at least 0.5 hr? Traffic flow theory suggests that in steady-state conditions such recoveries take place over extremely short time periods. There is a strong possibility that the unexpected duration arises because of the averaging technique used in the analysis. That possibility will be examined in the next subsection.

The second question about the speed increase relates to the flow rates at the start (8:20 a.m.) and end (8:52 a.m.) of it. Figure 4 suggests that flow rates after the transition are lower than those at the start of transition. One line of reasoning from steady-state conditions (i.e., no demand changes upstream or downstream of the data acquisition location) suggests that recovery takes place at almost identical flow rates in the congested and uncongested regions and that, therefore, the result in Figure 4 is a consequence of the averaging. A second line of reasoning suggests that these are not steady-state conditions and that the queue clears because of a demand decrease on the main line (QEW). Which of these two is correct is examined further in the next subsection.

In the shoulder lane, the plot is quite different. Where the other two lanes show speed and flow decreasing from 6:50 a.m. to 7:30 a.m., the shoulder lane has flow increasing as speeds decrease. The heavy line in Figure 5(A) suggests that operations were close to capacity and were not forced downward along the lower branch [i.e., capacity is about 1,400 passenger car units (pcu) per hour]. However, it appears more reasonable that there is some operation on both the upper and lower branches, with the necessary transitions between them. Consequently, there could be a higher capacity, approximately 1,700 pcu/hr, as shown in Figure 5(B). If this is the case, why again do the speed changes take so

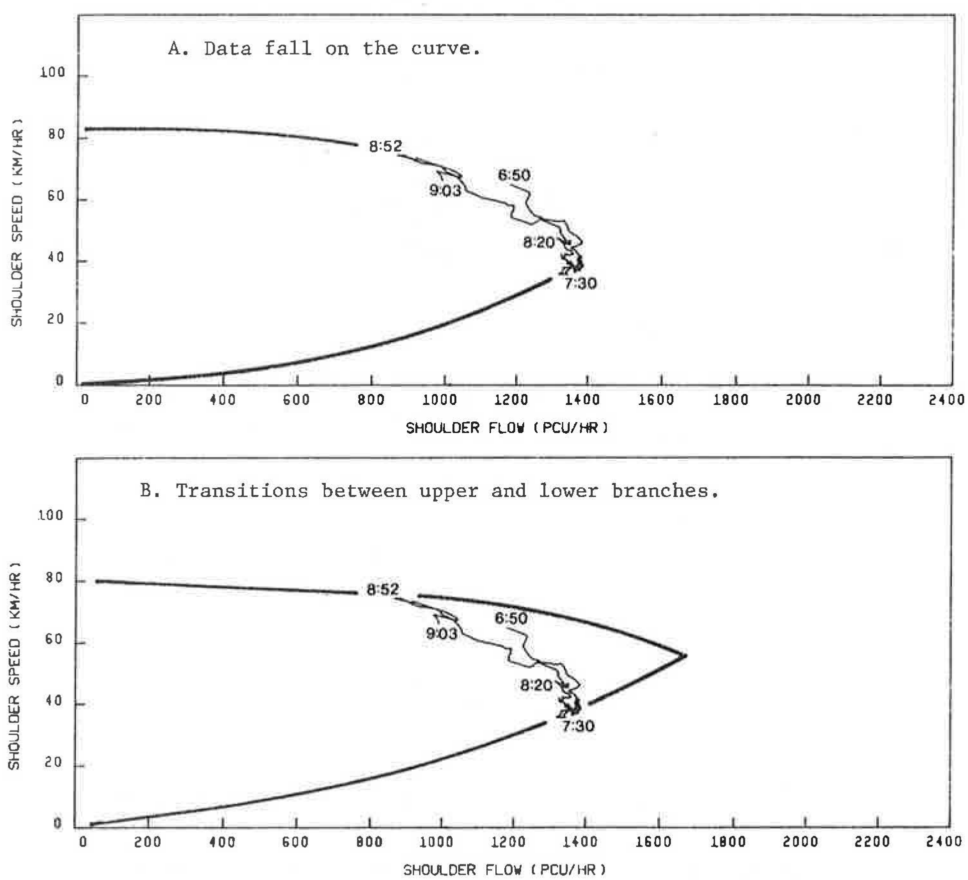

FIGURE 5 Representation of a speed-flow curve to fit shoulder lane data for Station 9.

long, with the end points at quite different flow rates? Particularly, during the supposed transition, why does the speed decrease with an increase in volume?

One possible answer is the same as suggested previously: the averaging technique may bias the view. Another possibility is that the station location strongly affects the result because it is immediately upstream of an entrance ramp feeding the bottleneck. To test this possibility, consider the plots in Figure 6 for Station 4, which is considerably upstream from any entrance ramps and at least 4 km from the bottleneck. In this instance, the operations on the shoulder lane, although at lower volumes, are similar in nature to those of the other two lanes with transitions to and from what appears to be definite lower-branch operation. It would, therefore, appear that the usual steady-state demand and shock wave notions are inadequate to describe the unusual speed-flow results for the shoulder lane at Station 9, which may be strongly related to site-specific driver behavior characteristics.

An important result of viewing the plots in Figures 4 and 6 is further confirmation of the approach of examining operations on a lane-by-lane basis instead of averaging across all lanes. When attempting to identify fundamental operating characteristics it is obviously important to avoid potential confusion and consider the lanes separately. In addition, it would appear equally important to examine more than one location, as indicated by the rather dramatic differences between shoulder lane operations shown in the two figures.

Event-Based Analysis

One possible explanation for the unexpected results in Figures 4 and 6 is that they arise because of the

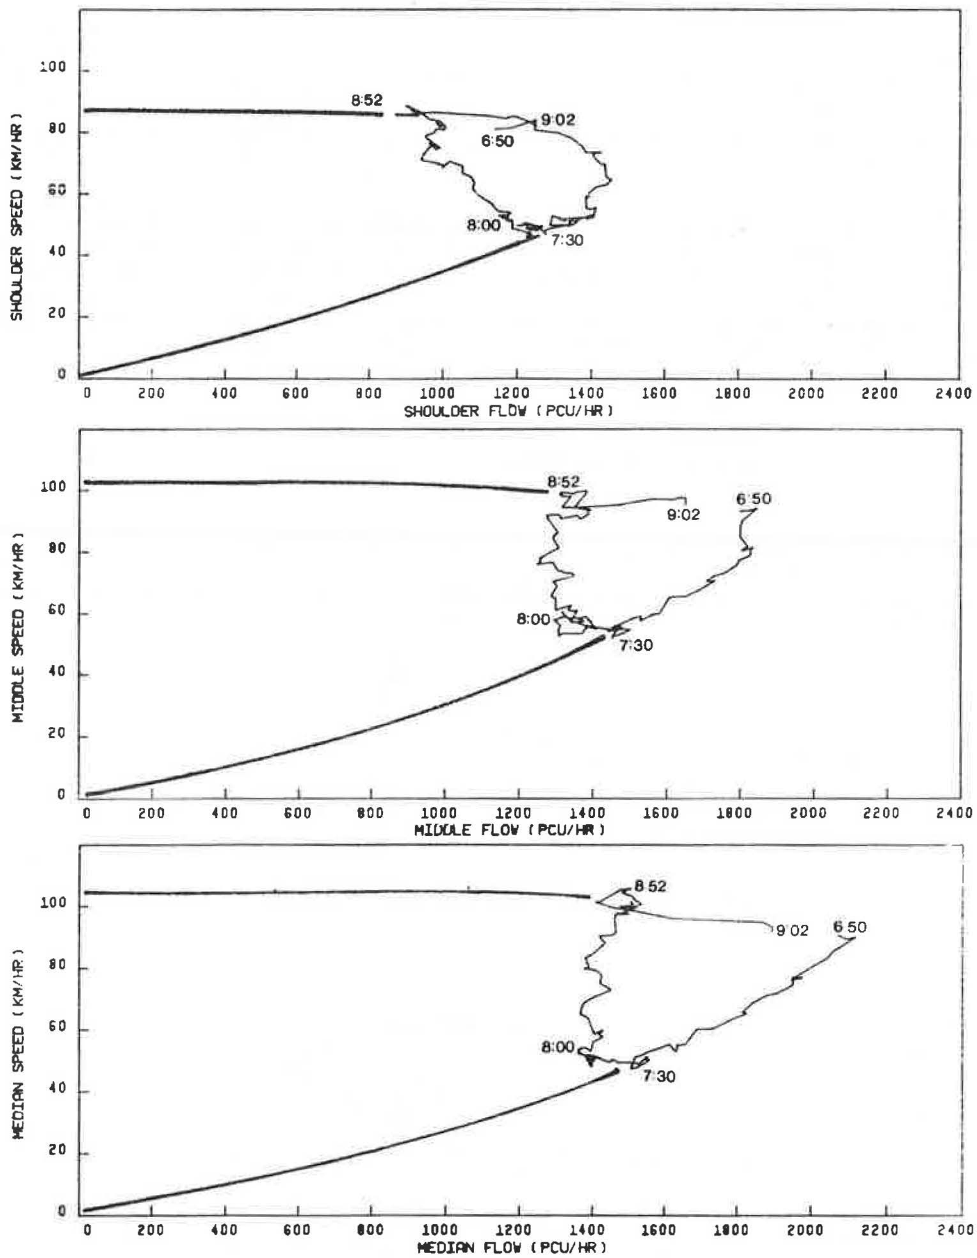

FIGURE 6 Station 4 mean speeds versus mean flow rates (over 68 days) connected by clock time; heavy lines suggest HCM-type curve consistent with the data.

use of moving average speeds and flows. In addition, averaging over clock time is not likely to mean averaging over identical traffic events. For example, congested operation will probably commence at slightly different times on different days. However, if the occurrence of a transition and average speeds and flows before and after that change in operation could be identified, without reference to clock time, a different and more representative picture might emerge. The procedures used to do that and the results for the transitions to and from congested (lower-branch) operation are described in this section.

The basis for identifying the change in conditions was the change in speed between successive intervals, because theoretically there will be a sudden decrease in average speed representing a move from the upper to the lower branch of the speed-flow curve, or the reverse. A change of 15 km/hr was selected on the basis that it was larger than most of the random fluctuations. In the case of the start of congestion, the first occurrence of a drop of 15 km/hr or more between successive 5-min average speeds (over the time period 6:50 a.m. to 8:00 a.m.) was taken to define the transition to congested flow. For termination of congestion, the last increase in speeds of more than 15 km/hr (between 7:30 a.m. and 9:10 a.m.) defined the recovery transition. This test was run for each day's data separately. Then the data were averaged across days on the basis of 5-min intervals before or after the identified speed shift. The results of these event-based calculations are shown in Figure 7 for Station 9 and in Figure 8 for Station 4. Both the increase and the

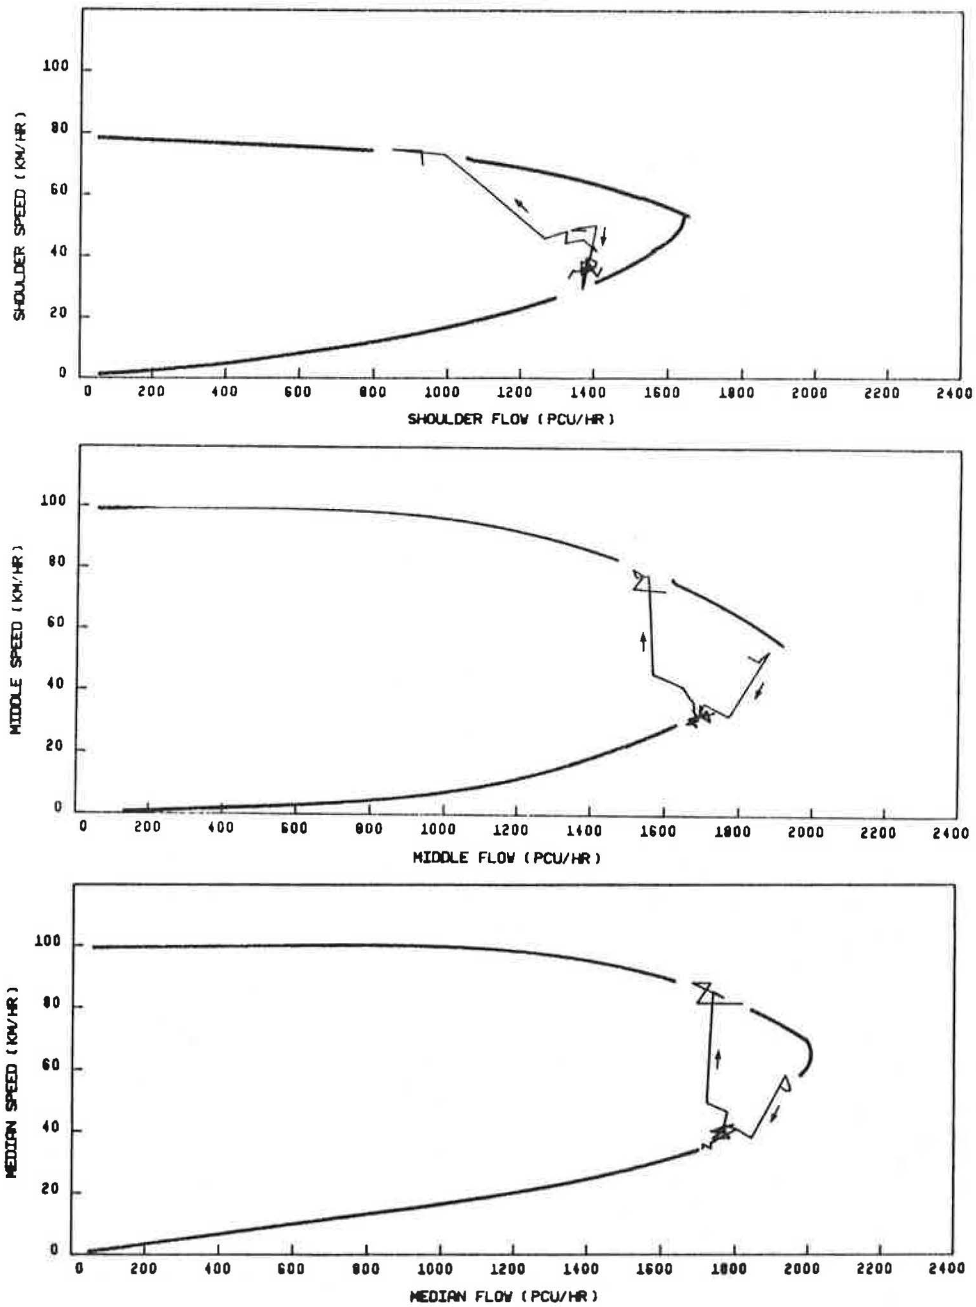

FIGURE 7 Station 9 event-based analysis of mean speeds and flows before and after transitions between congested and uncongested flows.

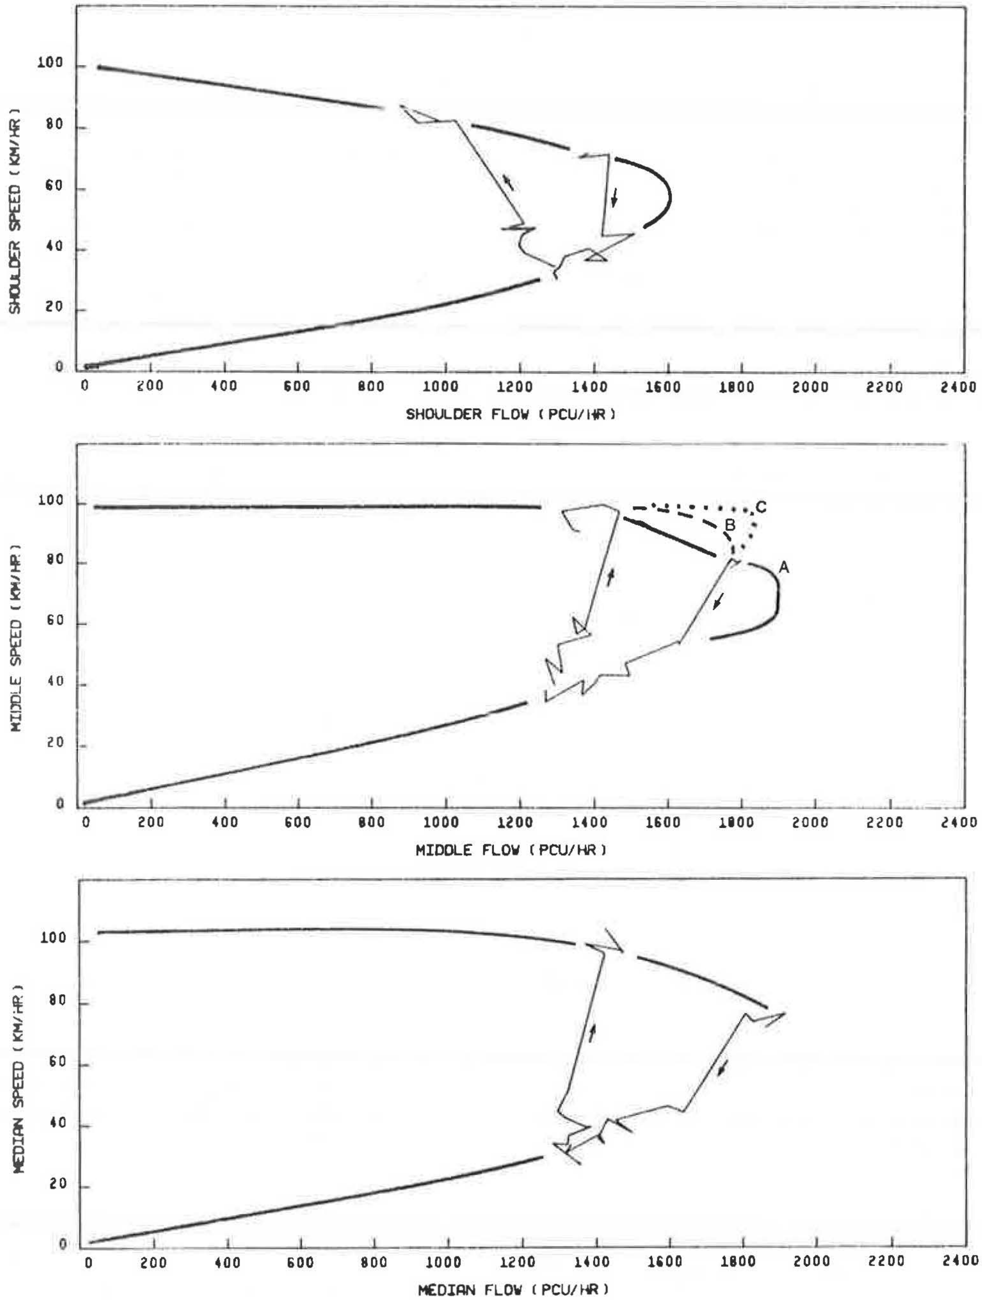

FIGURE 8 Station 4 event-based analysis of mean speeds and flows before and after transitions between congested and uncongested flow; for the middle lane, Line A represents HCM-type operation, Line B shows the simplest curve to put the data all on the lower branch, and Line C is suggested by Figure 6.

decrease are shown on the same figure to facilitate discussion, but each was calculated independently of the other.

These figures appear to resolve two of the unexpected results of Figures 4 and 6. First, both the increase and the decrease in speeds occur extremely rapidly. Within one 5-min interval, speeds have increased by nearly 40 km/hr (middle and median lanes, Figure 8), and in all cases by at least 25 km/hr. The speed decreases are not so large at Station 9, averaging only 20 km/hr, but the drop is still much quicker than was suggested by the time-traces in Figures 4 and 6. Thus conventional wisdom, which has the changes in condition occurring rapidly, appears to be supported.

Second, the increases in speed now appear to occur at roughly the same flow rate (shown most clearly in Figure 7, middle and median lanes) instead of with decreasing flow rates as shown in Figures 4 and 6. This is not so clearly the case at Station 4, or in the shoulder lanes, where the return to the upper branch is less direct. The results for Station 4 show flow rates increasing during the transition (for the middle and median lanes). Such a result suggests either that this analytical approach has not yet resolved all the difficulties or that the timing of main-line demand does not permit a consistent recovery in operations. To check the validity of the result shown in Figure 7, individual plots of daily time-connected speed-flow curves were examined to determine what kinds of recoveries actually occurred. In general (in the middle and median lanes), the large increase in speed occurred at roughly similar volumes, with some having slightly greater and some slightly lower volumes after the shift. The vertical lines in Figure 7, therefore, appear to be good representations of the average tendency.

Thus this event-based analysis appears to have resolved the two questions raised by the first analysis and has reinforced conventional understanding in several respects. However, one aspect of the results in Figures 7 and 8 is still unexpected because the downward shift is not located where it would be expected in two important respects: it starts from a lower speed than anticipated, and it is accompanied by a decrease in flow rates. It is also noteworthy that the transition ends at lower speeds (35-40 km/hr) than are suggested by the speed-flow curves of Figures 4 and 6 (45 km/hr). This last point confirms the potential hazards of averaging with respect to clock time, which in Figures 4 and 6 probably mixes congested and uncongested flow rates between 6:50 a.m. and 7:20 a.m., resulting in unrepresentative averages.

Three possible explanations for these differences can be suggested. The first is that the transitions observed occur between two points that are already on the lower branch of the curve as discussed earlier. This is shown in Figure 7 for the middle and median lanes. The resultant implied capacity (1,900-2,000 pcu/hr) is within the normally accepted range, and it is plausible that this particular location experiences capacity flows at or before 7:00 a.m. each day. If this is the explanation, then there is no surprise in these results.

Unfortunately, this explanation is not plausible for the Station 4 data (Figure 8). For one thing, observation suggests that Station 4 is not experiencing capacity flows by 7:00 a.m. Equally important, for the transitions to be wholly on the lower branch, it is necessary to accept average speeds of 80 km/hr as representing congested flow operation (for the middle and median lanes), which in itself contradicts conventional wisdom. Taking the middle lane as an example, there are three possible ways to sketch a speed-flow curve, each of which creates a problem. Line A in Figure 8 represents the conventional curve but leads to a problem in interpretation: there is a significant decrease in volume accompanying the transition from upper- to lower-branch operation. Line B is close to the conventional wisdom, but an attempt to put the start of the transition on the lower branch results in an abnormally low capacity for the section. Such low capacity appears to be unlikely. Line C is suggested by Figure 6, which shows quite high speeds at high volumes (which logically cannot be a false result of the averaging procedure). The problem with this representation is that the speeds at the start of the transition (85 km/hr) are then considerably lower than would be expected (from Figure 8, 100 km/hr).

These difficulties lead to the third explanation: this event-based averaging is still averaging dissimilar events and, therefore, producing misleading results. For example, on one day, flow just before the transition event may be 1,800 pcu/hr, and on another it may be only 1,600 pcu/hr. Close inspection of daily data may be the only way to resolve the problems. For now, there are three reasonable explanations for the observed results, none of which is fully satisfactory.

CONCLUSIONS

Three types of conclusions can be drawn from this work. The first relates to analytical procedures, the second to speed-flow relationships, and the third to further work.

The analytical procedures used here differ in four important respects from those of most previous work. First, the analysis has been conducted on a lane-by-lane basis instead of for the freeway as a whole. Mahabir's work first showed the importance of such an approach; it is believed that the current analysis provides further support for this approach.

Second, the analysis has not proceeded by the normal curve-fitting approaches, such as regression analysis. These are the obvious approaches to interpreting a large data set, but such curve fitting is inappropriate in the present context for two reasons. Fluctuations in the data make it impossible to tell from normal scatterplots what is upper-branch and what is lower-branch operation, with the result that the curve may well be fitted to inappropriate data. Also, if the focus of the analysis is on changes in operation (i.e., transitions from one branch to another) the "lines" of interest are not those represented by the average curve.

The third difference in the analytical procedures is that an event-based averaging procedure has been used instead of clock-time averaging in order to isolate the changes in operations. This procedure is not the final answer, however, as evidenced by the confusing results for the middle lane shown in Figure 8, but it does provide better answers than do other procedures about what actually occurs on freeways. Additional efforts are needed to develop more appropriate analytical methods for such data.

The fourth difference is the size of the data set. For an analysis of the type conducted here, data for a few days or a week would be inadequate. Examination of the daily speed-flow traces made it obvious that there is considerable fluctuation in operations both within a single day and between days, such that results based on only a few days' data should be interpreted cautiously. For normal curve-fitting approaches, a week's data of 5-min flows would appear to be sufficient, but for identifying the nature of shifts in operation, considerably more days of data are required. The size of the data set used here permits reasonable confidence in the trends that have been observed.

The second type of conclusion pertains to speed-flow relationships. Some of the present conclusions support conventional wisdom; others are simply not addressed by published results and probably warrant more investigation. There are three main conclusions in support of current understanding:

1. Both the breakdown and the recovery of speeds occur fairly rapidly, with major changes occurring during a 5-min interval;
2. Recovery from lower-branch to upper-branch operation appears to take place at an approximately constant volume; and
3. Different speed-flow relationships are needed to describe operations at different points along a highway (e.g., those sketched in for Stations 9 and 4 in Figures 4 and 6).

Two main observations have not been addressed, at least in published versions of the conventional understanding:

1. The location of the line representing transitions from upper- to lower-branch operation is not easily explained and
2. There are distinctly different speed-flow relationships on the different lanes, even well upstream of any entrance ramp (e.g., Station 4).

The third type of conclusion focuses on future work. These results come from only one section of highway, with a particular pattern of demand over time on the mainstream and on the entrance ramps, and with a pattern already affected by the freeway control system. Some of the results can be explained plausibly on the basis of this specific demand pattern. Work in other places is necessary to see if those explanations are sound. Alternatively, more detail on flow rates entering the section could be used to test the proposed explanations. In particular, the following points require investigation:

1. The recovery in speeds takes place at a lower flow rate than that at which the drop occurred. It is clear that demand flow rates are not constant over the hour for either the mainstream or the entrance ramps: a steady-state analysis is inappropriate. If flow rates on the entrance ramp downstream of Station 9 increase after Station 9 has begun lower-branch operation, then the observed flows at Station 9 must decrease. If subsequently mainstream demand decreases (because of smart drivers who time their trips well, for example), then the recovery to upper-branch operation will be from these reduced flows. If this is a plausible explanation, there may well be other expressway systems for which recovery flows are equal to or greater than flows at the onset of congestion.

2. The HCM-type speed-flow curves sketched in on many of the figures have been shaped to fit the present results, consistent with current theory, and in particular with current ideas of plausible values for capacity flow. If capacity were much higher (such as 2,400 or 2,500 pcu/hr), some of the interpretation would change. There have been recent analyses that suggest that capacity is at least 2,200 pcu/hr (4,12), but those analyses were based on conventional curve-fitting approaches and may well estimate capacity incorrectly. More work needs to be done using different analytical techniques to identify capacity flow rates.

3. Density (or lane occupancy) data have been ignored in this analysis in order to maintain a consistent discussion. These data also need to be incorporated using the types of techniques tried here instead of conventional curve fitting.

4. The present analysis has looked at each station in isolation, but clearly these data come from a system in which observed main-line flow rates are a function of downstream ramp volume (which is determined by the freeway control system) as well as of main-line demand. A complete understanding of the system requires information on how each of these volumes varies. Such an analysis would also help to make clear to what extent these results are or are not representative of freeway operation generally.

Overall, this analysis has contributed several useful points to understanding speed-flow relationships on freeways as they move to and from congested conditions. Further work is definitely needed to clarify several points. What appeared at first to be an abundance of data is in the end inadequate to resolve all the questions raised. Some insights have been gained. More are needed if the dissatisfactions with current depictions of freeway operating characteristics are to be resolved.

ACKNOWLEDGMENTS

The analysis described in this paper would not have been possible without the extensive efforts of Geddes Mahabir in making the data set manageable on the McMaster computer facilities several years ago. The assistance of the Ontario Ministry of Transportation and Communications in providing the data and their continued interest are also appreciated. Financial support came from the Natural Sciences and Engineering Research Council, through their Undergraduate Summer Research Award program, and from McMaster University. The authors would also like to acknowledge the helpful comments of two of the reviewers.

REFERENCES

1. Highway Capacity Manual 1965. Special Report 87, HRB, National Research Council, Washington, D.C., 1965, 411 pp.
2. R.P. Roess, E. Linzer, W.R. McShane, and L.J. Pignataro. Freeway Capacity Procedures. In Transportation Research Circular 212: Interim Materials on Highway Capacity, TRB, National Research Council, Washington, D.C., 1980, pp. 151-266.
3. B.L. Allen and S.M. Easa. Toward Improved Capacity and Level of Service Procedures for Ontario Highways. Final Report 2312-7814. Traffic Research Group, McMaster University, Hamilton, Ontario, Canada, 1978.
4. B.L. Allen and G.P. Mahabir. Speed/Flow and Capacity Relationships on Multilane Freeways--A Summary Report. Summary Report 2312-8311. Traffic Research Group, McMaster University, Hamilton, Ontario, Canada, Feb. 1983.
5. Highway Capacity Manual. Bureau of Public Roads, U.S. Department of Commerce, 1950.
6. R.P. Roess. Freeway Capacity Analysis Procedures: Literature Review and Annotated Bibliography. Office of Research and Development, FHWA, U.S. Department of Transportation, 1977.
7. M. Koshi et al. Some Findings and an Overview on Vehicular Flow Characteristics. Proc., 8th International Symposium on Transportation and Traffic Theory, 1981, V.F. Hurdle, E. Hauer, and G.N. Steuart eds., University of Toronto Press, Toronto, Ontario, Canada, 1983.
8. V.F. Hurdle and P.K. Datta. Speeds and Flows on an Urban Freeway: Some Measures and a Hypothesis. In Transportation Research Record 905, TRB, National Research Council, Washington, D.C., 1983, pp. 127-137.
9. S.M. Easa and A.D. May. Generalized Procedure for Estimating Single- and Two-Regime Traffic Flow Models. In Transportation Research Record 772, TRB, National Research Council, Washington, D.C., 1980, pp. 24-37.
10. A. Cedar. The Accuracy of Traffic Flow Models: A Review and Preliminary Investigation. Traffic Engineering and Control, Vol. 19, No. 12, Dec. 1978.
11. S.M. Easa. Selecting Two-Regime Traffic-Flow Models. In Transportation Research Record 869, TRB, National Research Council, Washington, D.C., 1982, pp. 25-36.
12. G.P. Mahabir. Speed, Flow, and Capacity Relations on Multilane Highways. M.Eng. thesis. McMaster University, Hamilton, Ontario, Canada, 1981.
13. E.R. Case and K.M. Williams. Queen Elizabeth Way Freeway Surveillance and Control System Demonstration Project. In Transportation Research Record 682, TRB, National Research Council, Washington, D.C., 1978, pp. 84-93.
14. D.P. Ryan and S.M. Breuning. Some Fundamental Relationships of Traffic Flow on A Freeway. Bull. 324. HRB, National Research Council, Washington, D.C., 1962, pp. 73-84.

Publication of this paper sponsored by Committee on Traffic Flow Theory and Characteristics.