Some New Data That Challenge Some Old Ideas About Speed-Flow Relationships

B. N. PERSAUD AND V. F. HURDLE

Several issues related to the upper branch of the speed-flow curve are addressed using data gathered at a freeway bottleneck in Toronto. Some important findings evolve, some of which challenge conventional beliefs. At low to moderate flows, speed on the upper branch is found to be insensitive to flow, a notion that is becoming increasingly accepted; at higher flows, the data suggest that speed decreases with increasing flow, but the fall-off is not nearly so precipitous as is commonly thought. The data further suggest that the presence or absence of an upstream queue is a more important variable than flow for predicting freeway speeds. Perhaps the most important finding is that belief in a precipitous speed drop may very well have resulted from a misinterpretation of data that arises because the speed of vehicles discharged from a queue varies with location in the bottleneck.

In this paper, questions relating to the upper, high-speed branch of speed-flow diagrams will be examined in some depth through the use of data collected on a freeway in Toronto, Canada. To an informed reader, it might be puzzling that such fundamental questions should still be of interest after so many years of research, but the authors believe they are not only deserving of study, but qualify as truly neglected areas. Furthermore, that neglect may very well lead to unjustified decisions regarding freeway construction or control, hence have serious implications for transport policy.

Most of the literature seems to take it for granted that there is a functional relationship between flow and speed and that this relationship can be represented by a curve with a well known and clearly defined shape. This paper also assumes that the relationship exists and can provide useful information for at least some purposes, but raises questions about whether the curve's shape is what most books would lead one to believe. The current state of the art in North America, as reflected in the speed-flow curves for freeways in the Transportation Research Board's 1985 *Highway Capacity Manual* (*1*), has evolved from numerous empirical studies, so one might suppose that little is to be gained by yet more empirical studies. However, the fact that these curves have gone through an evolutionary process and are quite different from previously used curves—those in earlier editions of the *Highway Capacity Manual*, for example—does not necessarily imply that they are correct, but should, instead, be taken as justification for further research that explores whether or not this process of evolution is complete. Further, in this research area, what constitutes current lore is often based on consensus of empirical research findings, so any contribution toward the formation of a consensus can be seen as worthwhile. Thus, even if the result of this type of research is merely a confirmation of current beliefs, a contribution will have been made in that one can have increased confidence in current beliefs. In the case of this paper, the results will tend to confirm some current beliefs, but give cause for doubting that others are correct.

Once one makes the basic assumption that it makes sense to talk about speed-flow curves at all, the questions about the upper branch that appear to be unresolved fall into two areas and can be illustrated using the curves from the *Highway Capacity Manual* (*1*) shown in figure 1. The first question concerns the part of the curve at low to moderate flows. Recent literature— Roess, McShane, and Pignataro (*2*), Hurdle and Datta (*3*), and Allen, Hall, and Gunter (*4*), for example—suggests that the upper branch is quite flat at these flows, at least for North American conditions where speed limits are set and enforced in such a way that rather few vehicles travel a great deal faster than the average speed. Indeed, the new *Highway Capacity Manual* states that "There is a substantial range of flow over which speed is relatively insensitive to flow; this range extends to fairly high flow rates." (*1, page 3-5*). However, despite this trend in current thinking, the bulk of the available literature suggests that speed drops even at quite low flows, so it seems likely that there are still some nonbelievers who need to be further convinced.

Furthermore, there is a real question as to just how insensitive to flow speeds are within the range of zero to perhaps 1,500 vehicles per hour per lane. Figure 1 shows a drop of 6 to 8 mph over this range, but the authors can see no evidence of such a drop in the data presented in the three works mentioned above (*2, 3, 4*). A comparison of the data points in

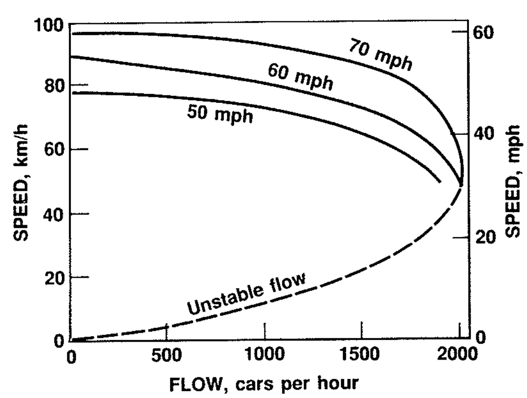

FIGURE 1 Speed-flow curves from the *Highway Capacity Manual* (*1*).

Department of Civil Engineering, University of Toronto, Toronto, Ontario M5S 1A4, Canada.

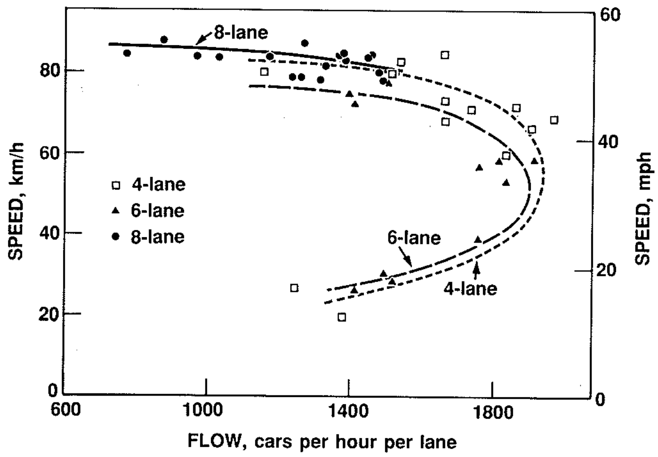

FIGURE 2 Speed-flow data and curves proposed by Roess et al. (*2*).

figure 2, taken from Roess, McShane, and Pignataro (*2*), and figure 3, taken from Hurdle and Datta (*3*) is particularly interesting in this regard. This is because the two figures also appear in the *Highway Capacity Manual* (*1*) and because there was a 55 mph speed limit on the New York area parkways where the figure 2 data were gathered, but a limit of 100 km/h (62 mph) at the Canadian location of the figure 3 study. This difference is clearly reflected in the two figures, but the authors can see no evidence in either of them that the average speed dropped as the flow increased from zero to 1500 vehicles per hour per lane. Certainly, it did not drop anything like the 6 to 8 mph indicated in figure 1.

The second major question area concerns the high-flow portion of the upper branch. The question here has three parts: At high flows, is speed no longer insensitive to flow? If it is not, where is the "break" point in flow, and what is the shape of the diagram at higher flows? Figure 1 suggests that the answer to this set of questions is that there is no real break, but that the slope of the curve changes very gradually at first, then increasingly rapidly. In the case of the 70 mph design speed, this slope change continues until the curve becomes vertical; in the words of a *Highway Capacity Manual* summary prepared for personnel of the Federal Highway Administration (*5*), ". . . speed precipitously declines as flow approaches capacity." While this answer appears to reflect common belief, some doubt lingers. Hurdle and Datta (*3*), for example, suggest that speeds on the entire upper branch are not a function of flows at all, but of whether or not vehicles have been in an upstream queue. The results of the current study do not support quite such an extreme view, but neither do they support the idea that there is a precipitous drop in speed as the roadway's capacity is approached. Furthermore, they indicate that the presence or absence of an upstream queue is a more important variable than flow if one wants to predict freeway speeds and that when there is a queue upstream, the speed is primarily a function of the distance from the observation point to the head of the queue.

Before presenting those results, however, it is instructive to review some of the existing literature in hopes of discovering the foundations of current thinking about the upper branch. This examination will provide a basis for judging current thinking and a backdrop against which results from this study can be presented. In reviewing previous empirical work, two bundles of issues—categorized according to whether they are conceptual or analytical—appear to be of primary interest. Along the way, some possible pitfalls will be discovered; it seems quite possible that some current beliefs about the upper branch may have come into being because of improper handling of these issues. In the current study, great care was taken to avoid these pitfalls, but the extent to which earlier studies avoided them is not clear from the published literature.

CONCEPTUAL AND ANALYTICAL ISSUES

The first issue to be examined arises from the authors' belief that the way one draws the upper branch is dictated by how one conceptualizes the entire speed-flow diagram. Probably the most common concept is that the upper branch represents free-flow conditions, the lower branch represents unsteady operation characteristic of conditions in a queue, and capacity occurs where the two branches meet. The authors have no quarrel with this concept, but believe that a problem arises

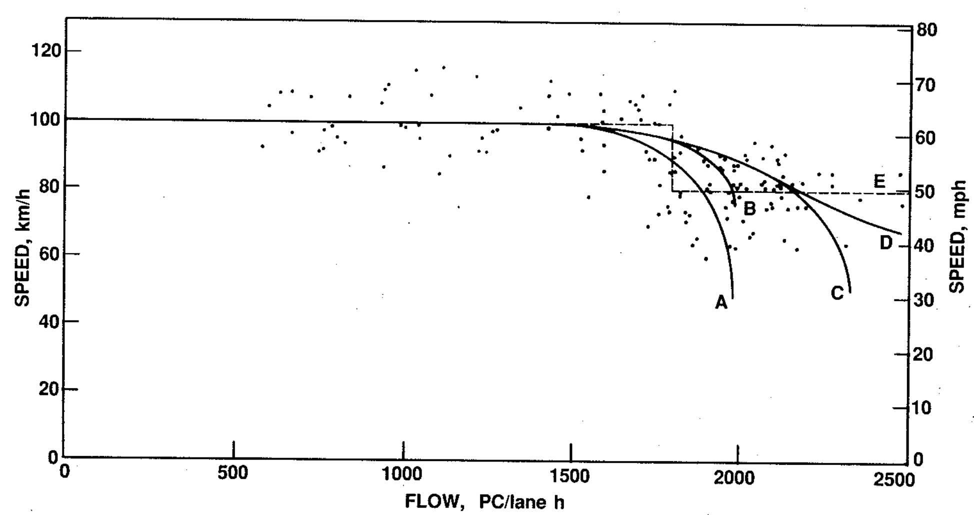

FIGURE 3 Speed-flow data and some possible curve shapes discussed by Hurdle and Datta (*3*).

because there has been little convincing evidence to indicate how the two branches meet, or if they meet at all.

It is very common to assume that the two branches form a single, smooth, continuously differentiable curve, but we can see no logical reason why this should be true necessarily. It is easy to see, however, that the shape of the upper branch one draws is very dependent on whether one assumes that the entire speed-flow curve is continuous, continuously differentiable, or neither. In particular, if one assumes continuous differentiability, then the curve must be vertical at the right end, as is the 70-mph curve in figure 1. Such curves are very common, but the authors suspect that most of them become vertical, not because of anything in the data on which they are supposedly based, but because some researcher had an a priori notion that the speed-flow curve—or, more likely, the speed-density curve—must be continuously differentiable. Probably the best-known empirical study addressing these issues is that of Drake, Schofer, and May (*6*), who fitted several of the well-known hypotheses—Greenburg, Underwood, Edie, Greenshields, and so on, to find out which one best fit their data. Some of these hypotheses imply a continuous speed-flow curve, some a continuously differentiable one, while some suggest a discontinuous curve of two or three regimes. After applying sophisticated statistical tests, the authors concluded that "the various hypotheses endured these tests with little differentiation." Unfortunately,this conclusion is probably primarily a result of the fact that statistical testing is a very blunt tool for the purpose. As discussed later in this section, Duncan (*7*), who relied more on subjective, visual methods and less on statistical procedures, instead concluded that some possible curves—in particular, those that are continuously differentiable—were not compatible with his data.

This paper's description of statistical testing as a blunt tool can be appreciated if one considers the problem of deciding which of the five speed flow curves in figure 3 best fits the data. There are two problems with using statistical testing procedures to solve this problem. The first is that the vertical scatter of the data points is so great that one would need a very large data set to show that the fit of one curve was significantly better than another, even when the curves differ as radically as those in figure 3. This difficulty is aggravated by the fact that data for high flow, uncongested conditions are difficult to obtain because such conditions ordinarily last for such a short time that only a few data points can be obtained from each day of observation. The second difficulty is one of definition: how does one statistically compare curves A and B, which become vertical at about 2,000 passenger car units per hour per lane, with curves D and E, which extend to the highest flows observed, but not to low speeds?

A third, closely related analytical issue has to do with curve fitting. If one wants to fit a curve to data by statistical methods, one must first specify an algebraic form for the curve, a parabola, for example, or some sort of logarithmic equation. The choice of algebraic form, however, can be more important in determining the shape of the curve than the data. Suppose, for example, that two researchers try to fit a curve to data similar to that in figure 2 by least squares procedures. Both decide to fit a curve of form

$$y = a + b_1x^2 + b_2x^2 + \dots,$$

but researcher A treats speed as the dependent variable (y) and excludes the observations at less than 30 mph as obviously not upper branch data, while researcher B treats flow as the dependent variable (y) and uses all of the data. Neither procedure can be criticized as obviously unreasonable, but for any data even vaguely resembling that in figure 2, researcher B will obtain a curve that becomes vertical at the right end: a result that researcher A cannot possibly obtain.

In presenting their data, the authors shall try to avoid this problem by not doing any numeric curve fitting, yet discussing the data as though a curve were being fitted. In essence, each reader is asked to think of fitting a curve to the data either by eye or by some marvelous, as yet uninvented, analytical method that can yield a curve which truly fits the data without the researcher having to specify a form for the equation. In doing so, any ability to produce numeric results in favor of a greater freedom to discuss the shapes of curves that exist only in the readers' minds is sacrificed. Naturally, the authors hope these curves will resemble those in their minds, but they have deliberately refrained from drawing any curves, preferring to rely on gentle persuasion rather than visual suggestion.

A second decision the authors made is not to show data points obtained when the study section was congested (*i.e.*, lower branch data). In part, this simply reflects the fact that their study section—described later in this paper—is a bottleneck, so normally causes congestion upstream rather than becoming congested itself. However, it does occasionally become congested, so there is a very limited amount of lower branch data. That it is not shown reflects the authors' doubts that the two branches of the speed-flow relationship form a single, smooth curve, but readers are asked to keep an open mind on this issue. This paper will also, initially, omit observations made while a queue existed upstream from the study section. That these observations do not constitute legitimate upper branch data is one of the paper's main points, but readers will eventually be shown the data and asked to judge for themselves.

The final issue in this bundle relates to the question of whether one examines the speed-flow relationship directly or indirectly. In tracing the evolution of speed-flow diagrams, one has to suspect that popular ideas about their shape arose from the work of investigators who first explored the speed density relationship, then inferred a speed-flow relationship from the expression flow = speed × density. Such an approach may well have seemed appropriate because speed and density were "natural" variables in car-following theory and because speed-density data has considerably less scatter than speed-flow data. Duncan (*7*), however, struck a telling blow against this philosophy in a landmark paper that, unfortunately, seems not to be as well known as it deserves to be. He first fitted two plausible relationships to some speed-density data, one continuous and one with a discontinuity, then calculated and plotted the corresponding speed-flow relationship for each. Next, he transformed the data to speed-flow form and,using some sound intuitive arguments, fitted new curves to this transformed data. Again, what was produced was a variety of shapes for the upper branch, but with the interesting feature that the shape of the curves based on the transformed data was radically different at the high flow end than the shape inferred from the fitted speed-density curves.

This issue can be further illustrated by examining it in the context of a data set and some curves presented by Leutzbach (*8*). Figure 4, prepared from figures in Leutzbach's paper, shows 1-minute speed-density data gathered at four locations

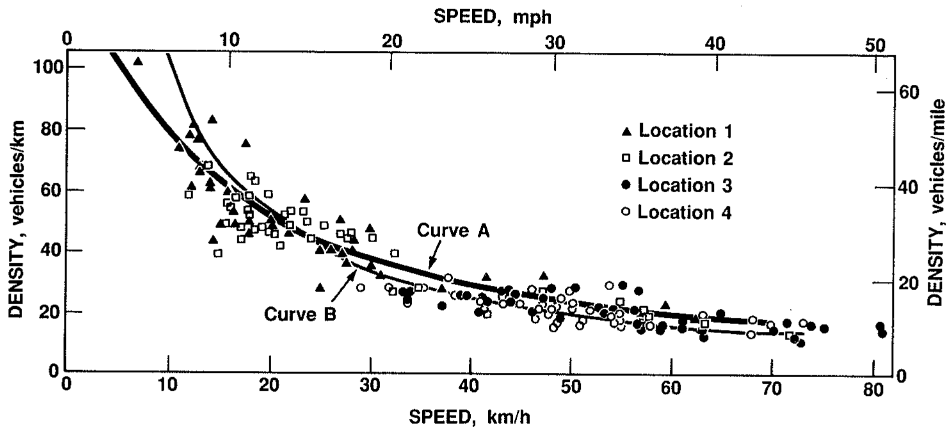

FIGURE 4 Speed-density curves prepared from figures in Leutzbach (*8*).

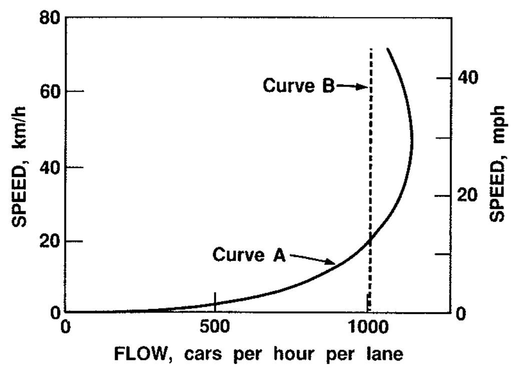

FIGURE 5 Speed-flow relationships implied by the curves in Figure 4.

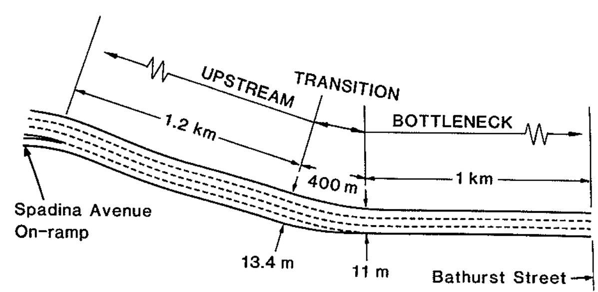

FIGURE 6 Plan view of study area (not to scale).

300 meters apart at the beginning of a one-lane bottleneck while there was a queue at the entrance. Leutzbach fit curve A by eye, and it appears to fit the data quite well; it translates into the speed flow relationship indicated by curve A of figure 5. Curve B of figure 4, also taken from a diagram in Leutzbach's paper, is a fitted speed-density relationship of the form speed = constant/density, implying, as one might expect of vehicles being discharged from a queue, that flow is reasonably constant in the bottleneck: the speed-flow relationship is merely a vertical line in figure 5. Thus, even though there is very little difference between the two speed-density curves in the region where data was available, the implied speed-flow relationships are radically different. From this illustration, there is once again a clear message: one must be cautious when transforming relationships from speed-density form to speed-flow form or vice versa. Specifically, one must be very cautious about using this type of transformation to form ideas about the shape of the speed-flow diagram.

EMPIRICAL RESULTS

The study (*9*) on which this paper is based was carried out in the vicinity of a bottleneck on the Gardiner Expressway in Toronto, Canada. As shown in figure 6, the Spadina Avenue entrance ramp joins the freeway and forms a fourth lane which is dropped within a 700 m radius curve after about 1.2 km to form a three-lane bottleneck. The speed limit is 90 km/h (56 mph) and the freeway has many restrictive design features, but speeds in excess of 100 km/h are common in the study area. Trucks are prohibited in the median lane, so flows expressed in vehicles per hour in this lane are, in effect, equivalent to flows in passenger car units per hour. Since this eliminates the complication of determining and applying passenger car equivalency conversions, most of the results presented in this paper will be for the median lane only.

Traffic leaving the downtown area during the afternoon rush period was observed by taking pictures with a 16 mm time-lapse camera mounted approximately 360 m above the ground on a tower located just off the left edge of the figure. The low to moderate flows in the opposite (inbound) direction were also captured on the film. For the outbound, high-flow direction, there were about 36 minutes of data over three days that showed free-flow conditions just before the queue formed upstream of the bottleneck and substantially more data after the queue had formed. These three days yielded about three hours of low to moderate inbound flow data as well. The outbound freeway segment was divided into sections as indicated in figure 6 and a 2-minute averaging interval was used for speed-flow measurements. With the data reduction method used (*9, 10*), average flows and densities were, in effect, obtained directly, while average speeds were computed from the expression flow = speed × density. A discussion of the accuracy of the method of computation and a comparison of calculated speeds with values obtained by direct measurement of individual cars' travel times is included in the authors' reference (*10*).

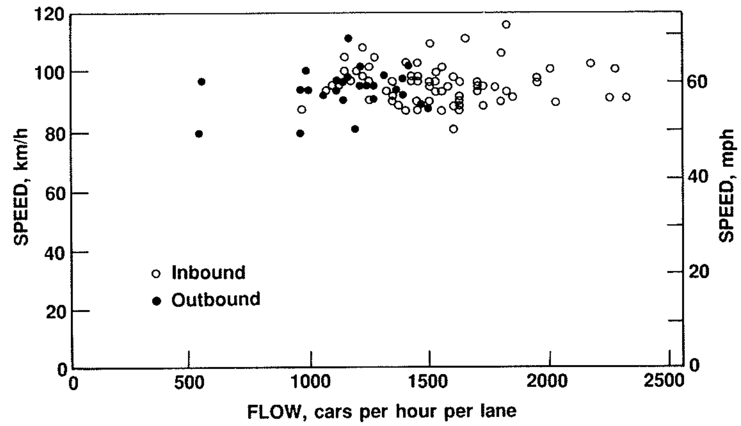

FIGURE 7 Speed-flow plot for inbound and outbound off-peak flows.

SHAPE OF THE UPPER BRANCH AT LOW TO MODERATE FLOWS

As indicated above, low to moderate flows prevailed in the inbound direction, so it was decided to examine the shape indicated by the inbound speed-flow data. The entire freeway bends in the area of the outbound lane drop and the inbound median lane segment for which data was extracted is just downstream of that curve. In figure 7, the speed-flow observations for this segment are indicated by open circles. To check whether the two directions have similar characteristics and to supplement the inbound data, some outbound median lane data were gathered during an off-peak period. These observations, which are denoted on figure 7 by the solid circles, indicate that speeds at moderate flows are about the same in both directions. The most striking aspect of figure 7, however, is that, for the flows of interest in this part of the paper—those up to about 1,800 vehicles per hour—speed in both directions fluctuates around an average of about 95 km/h at all flow levels, with no indication of a decrease in speed as the flow increases.

The final thing to note about figure 7 is that there are eleven inbound data points at flows larger than 1,800 vehicles per hour, all at speeds greater than 90 km/h, apparently indicating that the insensitivity of speed to flow applies to all flows. However, since these data points are few and scattered, they can be more properly discussed in the context of the relatively large number of high-flow outbound observations. This is done in the next two sections of this paper.

SHAPE OF THE UPPER BRANCH AT HIGH FLOWS

This section will explore the shape of the high-flow portion of the upper branch of the speed-flow curve. To do so, it is useful to first look at the outbound data on those days when there were observations without a queue present. As indicated earlier, three days of filming captured short periods of time when the outbound flows in the median lane were in excess of 1,500 vehicles per hour, and there was no noticeable queue. The three days' speed flow observations in the median lane during this period are plotted as x's in figure 8 and the outbound off-peak data points introduced in figure 7 as solid circles.

On each day, a queue formed upstream, so one can be sure that the capacity of the lane was reached. The x's include all 2-minute observations made before this happened, but none made after it occurred. Thus, whether they include conditions at capacity is perhaps questionable, but they certainly include conditions approaching capacity. (Readers are cautioned not to infer from figure 8 that the capacity of this lane is more than 2,400 vehicles per hour. The flows shown are based on counts only 2 minutes long, so the amount of random fluctuation is considerable. The highest flows observed would

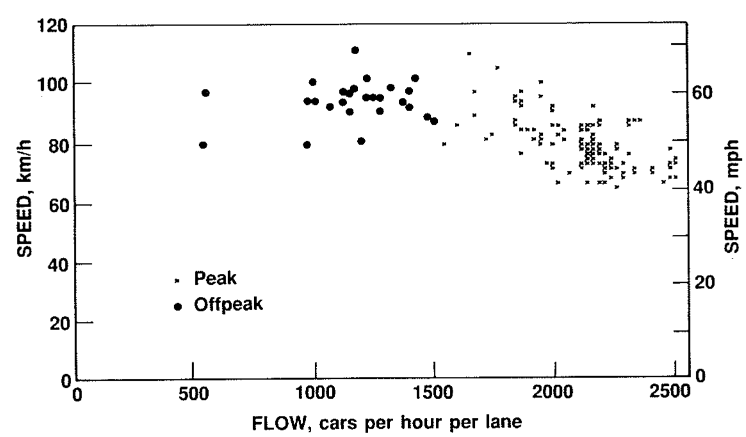

FIGURE 8 Speed-flow plot for the median lane, outbound.

undoubtedly be exceeded if one were to watch for a larger number of days, but they are never sustained for long periods of time. It would be unreasonable to say that just because they happened to occur, the capacity must be still higher.)

Visual inspection of figure 8 strongly suggests that, for flows exceeding 1,800 vehicles per hour, speeds do fall off gradually with increasing flows—quite a contrast to the pattern for lower flows. It is also apparent that the fall-off is not nearly so precipitous as the curves in figure 1 suggest. In fact, if a curve were to be fit by eye to the data for flows greater than 1800 vehicles per hour, a straight line would seem to appear. Even a curve with a small, but positive, second derivative seems compatible with the data, but the authors do not wish to suggest anything so radical. What they do want to suggest, however, is that a curve with a large negative second derivative such as the 70 mph curves in figure 1 seems incompatible with the data: if the curve had the shape shown in figure 1, at least some of the points at the right end of figure 8 would be expected to have lower speeds.

It would be tempting to try to estimate the shape of the speed drop and perhaps learn some more from the data by curve fitting or parameter estimation. As discussed earlier, however, the nature of the data provides a major stumbling block: the scatter in speeds is too large compared to the apparent change in the mean speed for these traditional techniques to be very useful without an extremely large data set. In addition, since data for many sections are combined, every vehicle is likely to be included in several observations; therefore, the data points are not all independent and statistical tests and the estimation of confidence limits on the parameters would be difficult, if not impossible, to carry out.

Figure 8 also suggests that the lowest speeds occurring at flows similar to, or in excess of, the flows normally quoted as capacity are of the order of 65–70 km/h (40–43 mph)—substantially higher than the "capacity" speed of 30 mph (48 km/h) indicated in figure 1. However, since figure 1 is based on average speed and flow per lane, the question arises: Do this and other findings based on median lane data apply to data averaged over all lanes? Because of the difficulty in accurately aiming the time-lapse camera, less data was available for the other lanes than for the median lane, but what is available (figure 9) clearly supports all of the conclusions so far about the shape of the speed flow curve. All that appears to be different is that average speeds and flows per lane are somewhat lower than in the median lane. The highest flow is now about 2,075 vehicles per hour per lane at speeds, as before, of 65–75 km/h (40 – 47 mph).

The finding that the fall-off in speed might not be so sharp as is commonly believed has some important implications. The first is that level-of-service criteria may need to be revised. Page 3-5 of the *Highway Capacity Manual* (*1*) states on the basis of the 70 mph curve in figure 1 that, "As capacity is approached, small changes in volume or rate of flow will produce extremely large changes in operating conditions, *i.e.*, speed and density. Level-of-service criteria for freeways reflect this, with the poorer levels defined for reasonably large ranges in speed and density, while the corresponding range in flow rates is rather small." If the fall-off in speed is neither as large nor as sudden as that shown in figure 1, then it is easy to see that predictions of the level of service to be expected at some given flow are likely to be overly pessimistic. The second implication is, in a sense, related to the first. Many believe that it is sound economic policy to prevent flows from reaching the levels at which the precipitous drop in speed occurs. Naturally, if the drop is not precipitous, such a policy would require rethinking.

Before concluding this section, it is worthwhile to give some special consideration to the interesting nature of the inbound flows in figures 7 and 9. Figure 10, which is a merger of the inbound and outbound observations from figures 7 and 8, clearly shows that, while for flows less than 1,800 vehicles per hour the two directions are visually indistinguishable, at larger flows there is little overlap. It is, therefore, tempting to suggest that the two directions behave differently at high flows, but with only eleven high-flow inbound data points, such a suggestion would require further support. However, one could speculate that under stable operating conditions a bunch of "brave" drivers can produce a high flow at high speeds in the median lane, but—as evidenced by the absence of high-flow, inbound data in figure 9—it is unlikely that such drivers would be found in all lanes during the same time interval.

POSSIBLE EXPLANATION FOR THE BELIEF IN A PRECIPITOUS SPEED DROP

Because the finding that there might not be a precipitous speed drop is so contrary to conventional wisdom and because

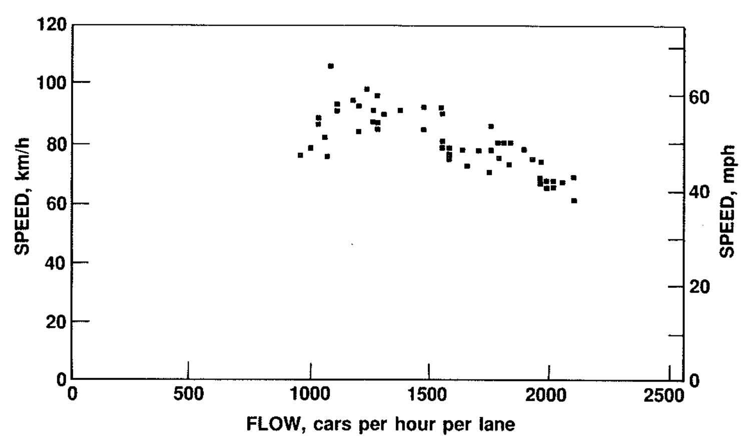

FIGURE 9 Speed-flow plot (3-lane average)—upper branch.

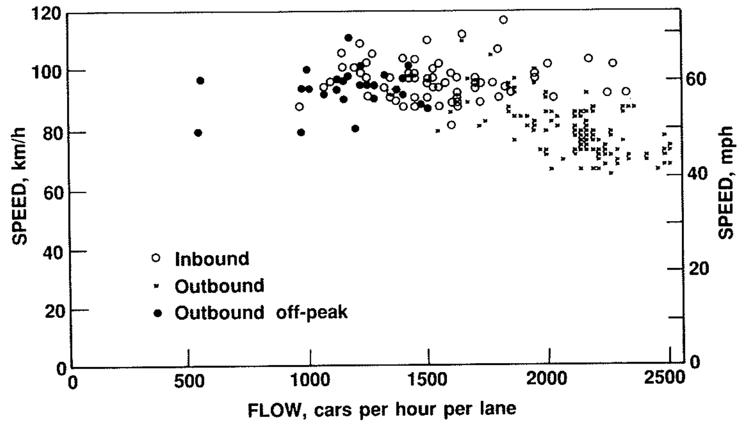

FIGURE 10 Speed-flow plot showing data from Figures 7 and 8 combined.

of the implications, it is natural to question whether such a finding from one empirical study can have wide applicability. If so, this would imply that studies on which current beliefs are based have erred in their conclusions. This is a difficult issue to address since it is not clear to what extent previous studies have sought to avoid the pitfalls hinted at earlier. It is possible, however, to point to at least one way in which one can erroneously arrive at a conclusion that there is a precipitous drop in speed at high flows.

The data points in figure 11, which is plotted with a different speed scale than the previous figures, are one day's speed-flow data observations averaged over all three lanes of the bottleneck while a queue was present upstream. Different symbols have been used to identify speed-flow data obtained in 110-meters long sections of the bottleneck; section 1 is located just past the lane drop in figure 6 and section 8 begins 768 m (not quite half a mile) farther downstream. For each section, the points appear to form a small cluster indicating reasonably steady speeds and flows. Average speed increases while average flow remains constant as one progresses into the bottleneck. The reason for this is that the vehicles upstream from the bottleneck are waiting in queue at either very low speeds or in the familiar stop and go fashion, so cannot be moving very fast in section 1 at the very upstream end of the bottleneck: simply because instantaneous speed change is not physically possible. Within the bottleneck, however, the vehicles do accelerate, so the speed gradually increases as one moves downstream. The only surprising thing about this is that the acceleration is so small and continues over such a long distance.

While the speed change data is of considerable interest in itself, the main reason for presenting it here is to point out that if one did not recognize the data points as "queue discharge" observations, they might easily be construed as supporting belief in a precipitous drop in speeds at high flows. It is easy to see this by combining the legitimate upper branch data in figure 10 with the "false" upper branch data in figure 11. The resulting plot, figure 12, clearly indicates how one can be led astray. This situation is not at all far-fetched; it is quite possible for upper branch speed-flow observations to "accidentally" become polluted by data such as that in figure 11, since data is usually gathered in such a way that it is difficult to tell if and when there is a queue upstream. In fact, very few published studies even say anything about whether there was a queue upstream. If data from more than one location are mixed together, as the authors suspect is often the case, the likelihood of misinterpretation becomes even greater since the different clusters of points, if all plotted with the same symbol, are likely to look exactly like the data one would expect to see if speeds dropped precipitously as the flow approached capacity.

SUMMARY

In this paper, several issues related to the upper branch of the speed-flow curve have been addressed using data gathered in the vicinity of a freeway bottleneck. Some important findings have resulted. There was confirmation of current belief that at low to moderate flows speed is insensitive to flow,

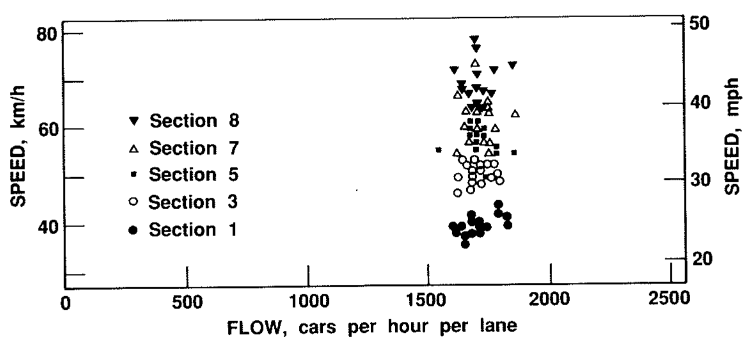

FIGURE 11 Speed-flow observations (3-lane average) at different locations in the bottleneck with a queue upstream.

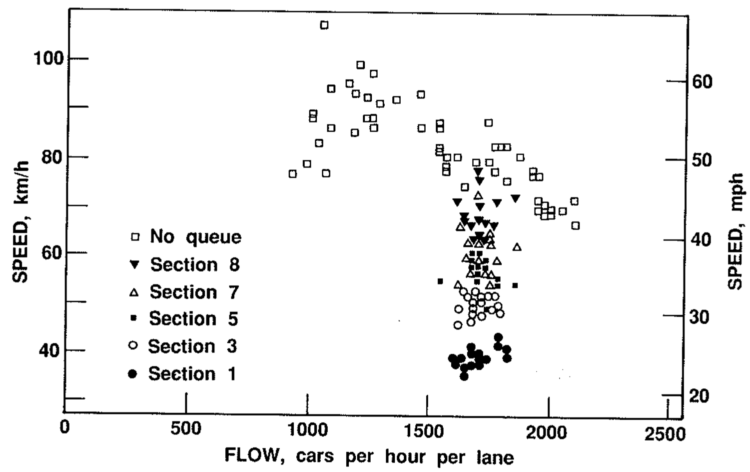

FIGURE 12 Merger of Figures 9 and 11 (at the scale of Figure 11) to illustrate "false" precipitous speed-drop.

though these results are in closer agreement with speed-flow curves in several of the references (*2, 3, 4*) than with the one in the *Highway Capacity Manual* (figure 1). At higher flows the study's data suggest that speed does decrease with increasing flow, but that the fall-off may not be nearly so precipitous as is commonly believed. Perhaps the most striking finding is that belief in a precipitous drop in speed at high flow may very well have resulted from misinterpretations of data that arose because the speed of vehicles discharged from a queue varies with location in the bottleneck, but the flow does not.

ACKNOWLEDGMENT

This research was supported by the National Science and Engineering Research Council of Canada.

REFERENCES

1. *Special Report 209: Highway Capacity Manual.* TRB, National Research Council, Washington, D.C., 1985, 516 pp.
2. R. P. Roess, W. R. McShane, and L. J. Pignataro. Freeway Level of Service: A Revised Approach. In *Transportation Research Record 699*, TRB, National Research Council, Washington, D.C., 1979, pp. 7–16.
3. V. F. Hurdle and P. K. Datta. Speeds and Flows on an Urban Freeway: Some Measurements and a Hypothesis. In *Transportation Research Record 905*, TRB, National Research Council, Washington, D.C., 1983, pp. 127–137.
4. B. L. Allen, F. L. Hall, and M. A. Gunter. Another Look at Identifying Speed-flow Relationships on Freeways. In *Transportation Research Record 1005*, TRB, National Research Council, Washington, D.C., 1985, pp. 54–64.
5. *The 1985 Highway Capacity Manual—A Summary.* Office of Traffic Operations, Federal Highway Administration, U.S. Department of Transportation, 1986.
6. J. S. Drake, J. L. Schofer, and A. D. May. A Statistical Analysis of Speed Density Hypotheses. In *Highway Research Record 154*, HRB, National Research Council, Washington, D.C., 1967, pp. 53–87.
7. N. C. Duncan. A Further Look at Speed/Flow/Concentration. *Traffic Engineering and Control*, Vol. 20, 1979, pp. 482–483.
8. W. Leutzbach. Testing the Applicability of the Theory of Continuity on Traffic Flow at Bottlenecks. *Proc.*, 3rd International Symposium on the Theory of Traffic Flow, New York, 1965 (L. C. Edie, R. Herman, and R. Rothery, eds.), New York, 1967.
9. B. N. Persaud. *Study of a Freeway Bottleneck to Explore Some Unresolved Traffic Flow Issues.* Ph.D. dissertation. University of Toronto, Toronto, Canada, 1986.
10. B. N. Persaud and V. F. Hurdle. Freeway Data from Time-lapse Film—an Alternative Approach. *Transportation Science*, in press.

Publication of this paper sponsored by Committee on Highway Capacity and Quality of Service.