UC Berkeley
Working Papers

Title
"Study of Traffic at a Freeway Merge and Roles for Ramp Metering"

Permalink
https://escholarship.org/uc/item/744926sv

Authors
Cassidy, Michael J.
Rudjanakanoknad, Jittichai

Publication Date
2002-05-01

eScholarship.org Powered by the California Digital Library
University of California

CALIFORNIA PATH PROGRAM
INSTITUTE OF TRANSPORTATION STUDIES
UNIVERSITY OF CALIFORNIA, BERKELEY

Study of Traffic at a Freeway Merge and Roles for Ramp Metering

Michael J. Cassidy
Jittichai Rudjanakanoknad
University of California, Berkeley

California PATH Working Paper
UCB-ITS-PWP-2002-2

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation; and the United States Department Transportation, Federal Highway Administration.

The contents of this report reflect the views of the authors who are responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the State of California. This report does not constitute a standard, specification, or regulation.

Report for MOU 3011

May 2002

ISSN 1055-1417

Study of Traffic at a Freeway Merge and Roles for Ramp Metering

By Michael J. Cassidy and Jittichai Rudjanakanoknad
University of California
Department of Civil and Environmental Engineering
and the Institute of Transportation Studies
109 McLaughlin Hall
Berkeley, California 94720
(510) 642-7702
Cassidy@ce.berkeley.edu

ABSTRACT

Traffic data measured near the junction of a single-lane on-ramp (with metered inflows) and a three-lane freeway were carefully studied for four days during the rush. The data showed the area around this merge junction became a bottleneck each day when the on-ramp's meter allowed its inflows to rise in the presence of high flows arriving from the freeway. Detailed study during these times further showed that queueing actually arose some distance downstream of the merge and that these queues were caused by drivers who, having just entered the freeway's shoulder lane from the on-ramp, slowed-down to negotiate discretionary lane-change maneuvers. The slowing was observed to spread quickly to all three freeway lanes and this had marked effects on the discharge flows from the bottleneck; i.e., average queue discharge rates were as much as 10 percent lower than the flows that had departed the merge prior to the bottleneck's activation. The latter of these flows, moreover, could be observed for many minutes (and therefore cannot be attributed to statistical variation).

Even after the bottleneck had activated, its discharge flows continued to be influenced by vehicle slowing at, or just downstream of, the merge; i.e., queue discharge rates temporarily diminished with each recurring slow-down. Reductions of 800 vph or more could accompany such a slow-down and these lower rates could persist for some minutes. These undesirable effects were again triggered by vehicles from the on-ramp changing lanes soon after having merged. And notably, the discharge reductions became more (less) pronounced as the ramp's inflows surged (diminished).

The findings show that service rates at this freeway site can be made higher 1) by eliminating, or at least postponing, the bottleneck's activation; and 2) by mitigating downstream slow-downs if or when the activation occurs. They further show, at times in dramatic fashion, that these objectives can be realized by carefully metering the on-ramp to limit the rates its vehicles join the freeway traffic stream. Certain issues concerning ramp metering in light of the present findings are discussed at the conclusion of this manuscript.

1. INTRODUCTION

This manuscript reports on certain traffic details measured at a freeway merge before and after it became a bottleneck each day. The findings indicate that flows departing this merge can be made substantially higher by metering the on-ramp to restrict the rate its vehicles proceed downstream. It follows that such a metering scheme can reduce the total delay at this merge (collectively incurred by all vehicles arriving from the freeway and the ramp), even absent any changes in travel patterns in response to the scheme.

These findings came to light by collecting high resolution data (including individual vehicle arrival times) at strategic locations near the merge and then processing and analyzing these data in careful ways. These painstaking procedures verified, among other things, that during each afternoon rush the site became an *active* bottleneck; i.e., queues formed downstream of the merge, but traffic remained freely flowing at locations just downstream of this (Daganzo, 1997). During its active periods, measured flows departing the bottleneck were therefore not impeded by any queues emanating from even further downstream.

The freeway site itself, and the kinds of traffic data extracted there, are described in section 2 of the manuscript. Presentations of the study findings are provided in the two sections that follow thereafter.

In section 3, data are displayed to reveal how on-ramp inflows affected certain traffic features around the merge. Specifically, these presentations demonstrate that the bottleneck activated at the downstream end of the merge when the on-ramp's inflows surged in the presence of high arrival rates from the freeway. They also show that the bottleneck's activation was marked by vehicle slowing that 1) originated in the shoulder lane just downstream of the merge; 2) spread quickly to the freeway's two other lanes; 3) periodically reoccurred with surges in the on-ramp's inflows; and 4) impeded discharge flows from the merge. The data further indicate that these periodic slow-downs were triggered by excessive numbers of drivers from the on-ramp who, having merged into the freeway shoulder lane, negotiated lane changes soon thereafter.

The observations displayed in section 4 verify that the merge area's service rates (i.e., its capacities) can be made higher with suitable on-ramp metering schemes. The presentations here rely upon measurements from a separate day when, shortly after the bottleneck's activation, the ramp's inflows were severely curtailed for several minutes. This abated the vehicle slowing just

downstream and produced a marked increase in discharge flows that persisted for some minutes.

Further, the observations from this (second) day show that eliminating or postponing the bottleneck's activation can also bring about higher service rates. On this day, the bottleneck activated only after a very high flow of nearly 7,000 vph had departed the merge. This high rate persisted for an extended time (15 mins) and was substantially higher than the ensuing queue discharge rates.

It is thus fortuitous that careful assessments of the data from all four-observation days confirm that metering more and more restrictively in the face of increasing arrival rates from the freeway can at least forestall bottleneck activations and prolong higher service rates. The data actually reveal a clear correspondence between the maximum measured rates departing the merge and ramp inflows, with lower values of the latter coinciding with higher values of the former. The evidence of this includes measurements from a day when the on-ramp's meter was not in service, such that surges in its inflows were especially high (and service rates, in turn, especially low). These presentations are also part of section 4.

The manuscript's fifth and final section summarizes some implications of these findings for choosing on-ramp metering schemes. Future experiments for the further testing of certain schemes are briefly discussed here as well.

2. STUDY SITE AND ITS DATA

Figure 1 is a sketch of the study site, a stretch of eastbound Freeway 22 and its junction with the Fairview Avenue on-ramp in Orange County, California. Loop detectors are located in each lane at MilePosts (MPs) 8.7 and 9.0. These recorded, among other things, vehicle counts over 30-sec intervals.[1] But as will be made clear in the following sections, data from these loop detectors played a minor role in the study. Instead, analyses were primarily performed on the individual vehicle arrival times at MPs 9.1, through 9.4 (as labeled on the figure). Vehicle trip times between some of these MPs were measured as well. These high-resolution data were manually extracted from videotapes that were recorded by cameras placed on the freeway over-crossing just downstream of the merge.

[1] During one of the days studied here (November 8, 2000), only the detectors at MP 8.7 functioned while those downstream at MP 9.0 did not record data. On the other three days, the opposite occurred; i.e., the detectors at MP 9.0 were the only ones to function.

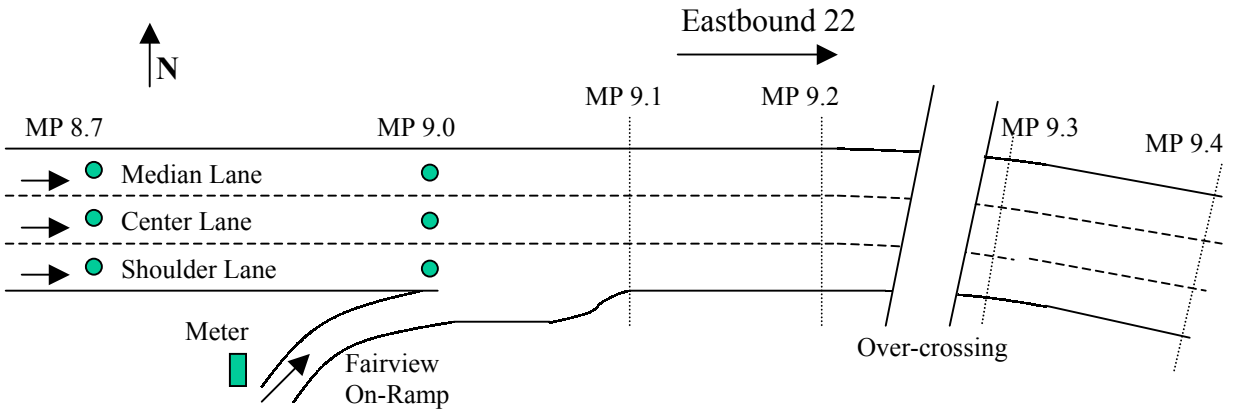


Figure 1
Study Site. Eastbound Freeway 22, Orange County, California

One of these cameras provided measurements of on-ramp vehicles passing MP 9.1 and merging onto the freeway. Notably, the meter at this ramp evidently varied its rates in response to traffic conditions measured (by the loop detectors) on the freeway. As a result, a wide range of on-ramp inflows was observed. Included here are some especially high surges that were measured on a day when the ramp meter was not in service.

In all, data were collected during the afternoon rush hours on four days spanning the period of November 8, 2000 to March 15, 2001. Weather conditions were good on each day and findings from the analyses of these data are provided in the following two sections of the manuscript.

3. SOME OBSERVED DETAILS OF TRAFFIC FLOW

This section reports findings from one of the four observation days (March 15, 2001). The presentation begins by verifying that the site became an active bottleneck. To this end, Figures 2(a) and (b) show curves of cumulative vehicle count, N, vs time, t, that were measured across all freeway lanes at certain MPs and transformed as follows. First, the curves in each figure were measured from the passage of an imaginary reference vehicle so that each set of curves was constructed from the same collection of vehicles. Secondly, each N-curve was shifted forward by the average free flow vehicle trip time from its respective MP to the downstream-most MP shown in its figure. The vertical separations between any two curves (of the same figure) are thus the excess vehicle accumulations between the respective MPs due to freeway traffic delays.

The N-curves were transformed in one final way to render these separations and certain other of their features more visible. Namely, an oblique coordinate system was used in each figure to plot

$N - q_o \cdot (t - t_o)$ vs t for each curve's starting time, t_o, and some choice of background flow, q_o; a fixed q_o was used for the entire collection of curves shown in a given figure and its value was selected so that the range of $N - q_o \cdot (t - t_o)$ was small as compared with the N itself. This coordinate system reduced the vehicle count actually displayed on each figure's ordinate. This, in turn, amplified the curves' vertical separations and made them more visible to the naked eye. The curves' changes in slopes were amplified as well, such that flow changes at a MP are made evident by taking piece-wise linear interpolations of the curves themselves and examples of this are provided in these and other figures.

Further discussion on the construction and interpretation of oblique N-curves is available from a number of sources, including Cassidy and Windover (1995). It suffices for now to direct the discussion to Figure 2(a). It presents curves at the four MPs downstream of the on-ramp during a certain 14-min period of the rush. It is clear from the figure that traffic was nearly freely flowing during the early minutes of this period; i.e., these curve portions remained nearly (although not entirely) superimposed, even as flows departing the merge climbed well above 7,000 vph. But by t = 13:30, the curves at MPs 9.1 and 9.2 diverged from their two downstream counterparts (indicating the presence of excess vehicle accumulations), while the curves at downstream MPs 9.3 and 9.4 remained superimposed. The freeway segment between MPs 9.2 and 9.3 was thus identified as the initial location of an active bottleneck; i.e., the head of a queue first resided somewhere between these MPs.

But Figure 2(a) also shows that queueing downstream of MP 9.1 all but disappeared by t = 13:34. By this time, its N-curves were again nearly superimposed and only occasionally displayed (rather small) vertical displacements.

Figure 2(b) shows, however, that the head of the queue had merely moved backward (against the flow of traffic) and that delays and excess vehicle accumulations persisted upstream of the merge. Shown in this second figure are oblique N-curves at MPs 9.0 and 9.1. These were constructed from the counts of vehicles that had arrived to the merge from the freeway; vehicles from the on-ramp were excluded here since these did not pass the location described by the upstream curve (i.e., at MP 9.0)[2]. The figure clearly indicates that queueing between MPs 9.0 and 9.1

[2] The N-curve at MP 9.0 was constructed from (30-sec) counts measured by loop detectors that were likely subject to small errors. But the finding that queueing persisted between MPs 9.0 and 9.1 rests upon large vertical separations between the respective N-curves, such that small errors in the detector counts would not change this

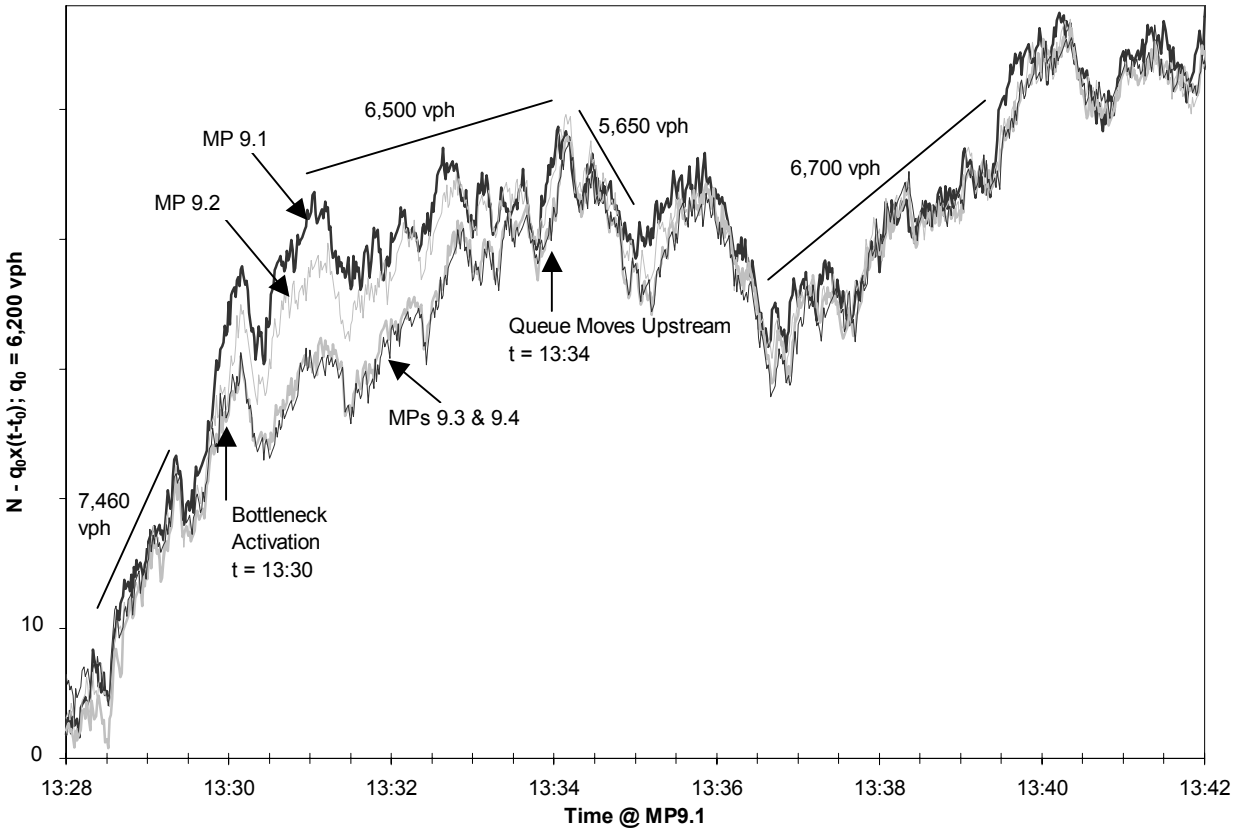


Figure 2(a)
Oblique N-curves at MPs 9.1 through 9.4 (March 15, 2001)

continued well beyond t = 13:34.

Figures 2(a) and (b) collectively verify that, beyond t = 13:30 (and prior to the bottleneck's de-activation some time later), flows measured downstream of the on-ramp (e.g. at MP 9.3) were queue discharge rates. These were the maximum rates that could depart the active bottleneck.

It is notable that, following the activation, the N-curves in Figure 2(a) display obvious changes in slopes, indicating changes in the queue discharge flows. These changes are evident both in the short-run (as wiggles on the curves) and in the longer-run (as exemplified by the linear approximations that accompany some extended portions of the N-curves). Longer-run trends in the curves also show that flows departing the merge diminished following the bottleneck's activation.

The explanations for these changes in flow begin with Figure 2(c). It displays vehicle trip times sampled over a stretch of the shoulder lane that lies between MPs 9.1 and 9.2 (since the

---

diagnosis.

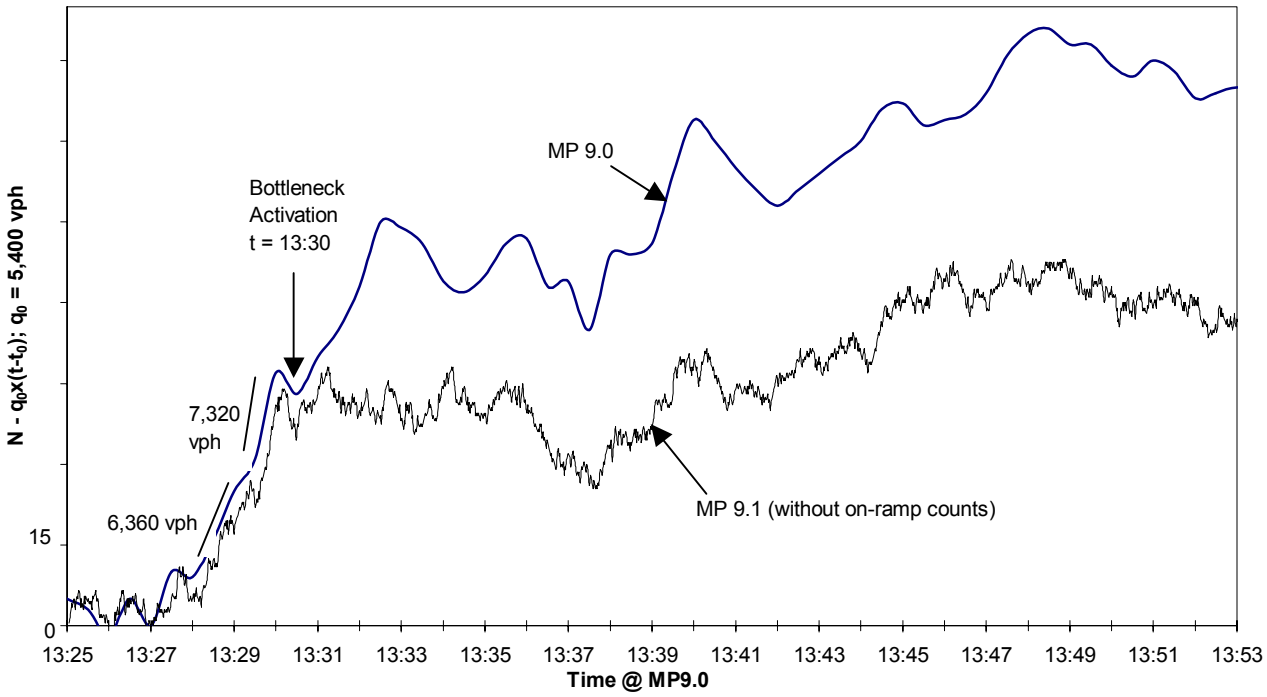


Figure 2(b)
Oblique N-curves of freeway traffic at MPs 9.0 and 9.1 (without on-ramps) (March 15, 2001)

videotaped images facilitated trip time measurements here). The figure displays the ratios of the actual trip times to the average free flow trip time; the latter was estimated from samples drawn early in the rush.[3] The data are expressed as these ratios to facilitate comparisons with other lanes and on other days (since the segment lengths used for sampling trip times changed in each of these instances). The ratios are plotted against the times vehicles arrived at the upstream end of the stretch and the inter-arrival times (i.e., the headways) between sampled vehicles were only about 5-secs long. The trip time ratios for each vehicle sampled are shown in the figure with the thin line. The bold line is the 30-sec moving average of these.

Figure 2(c) shows that trip times began to rise in the shoulder lane by t = 13:30, the time previously identified as marking the bottleneck's activation. They rose sharply soon thereafter and then diminished by t = 13:34, the time when the head of the queue had moved upstream.

[3] Designating the denominator of this ratio as the free flow trip time may be a misnomer; i.e., the denominator may actually be slightly higher than the average trip time under very low flows. This detail would not, however, affect our diagnoses.

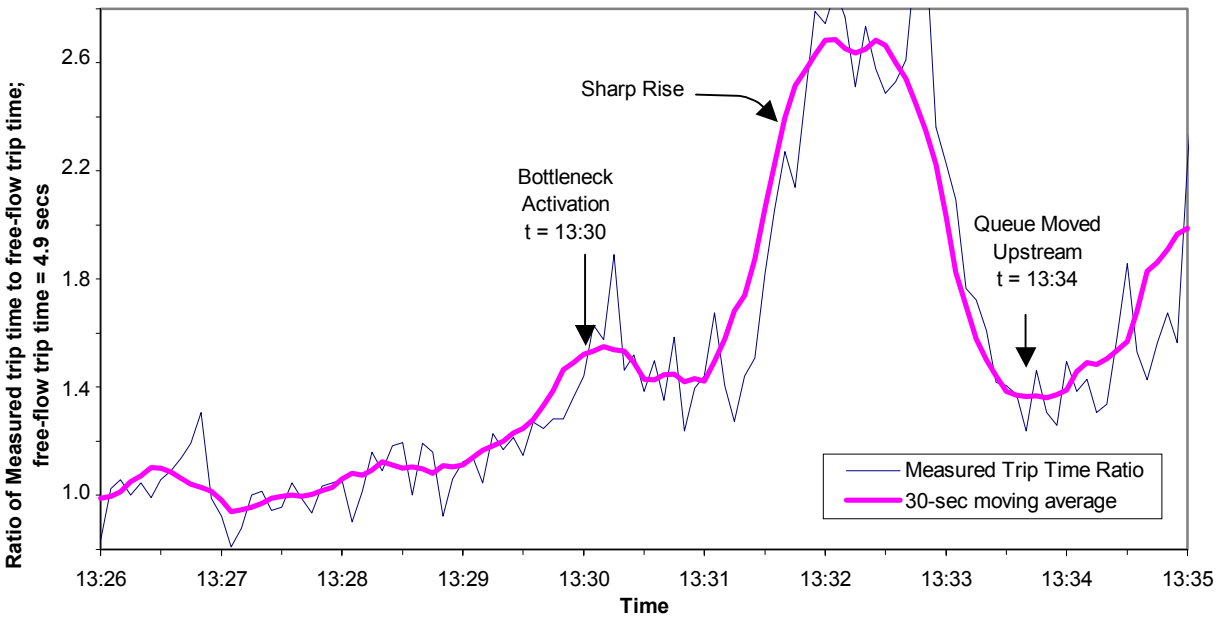


Figure 2(c)
Sample Trip Times in the shoulder lane (March 15, 2001)

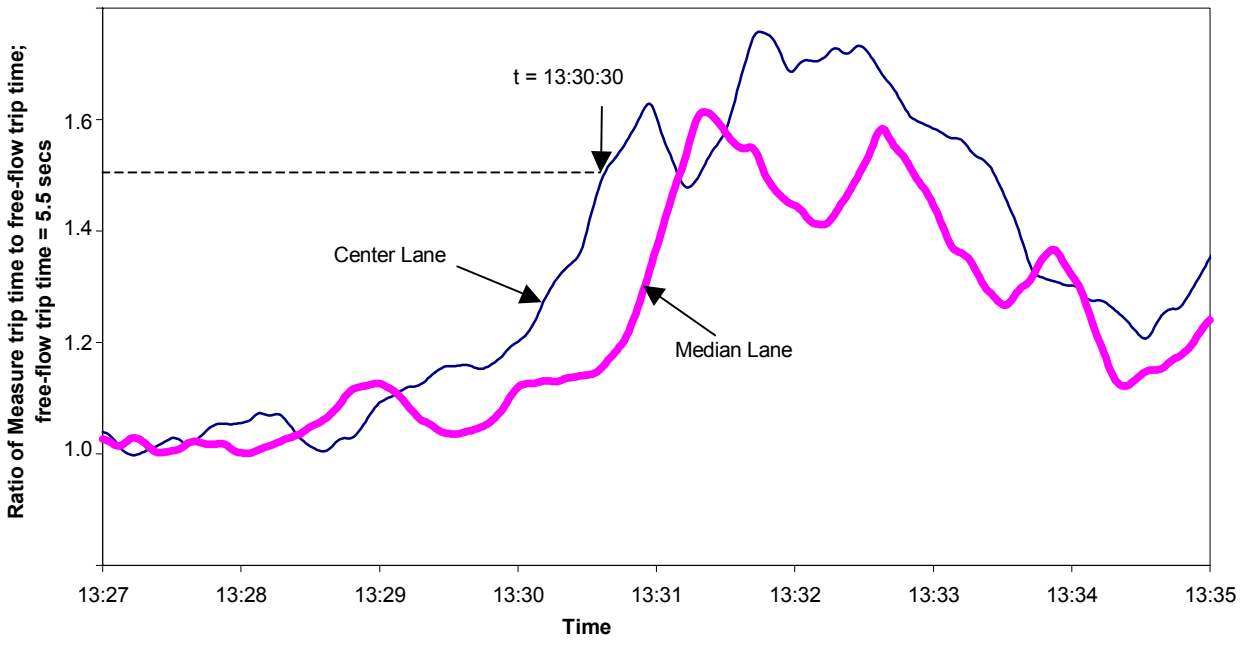


Figure 2(d)
Sample Trip Times in the center and median lanes (March 15, 2001)

Slow-downs in the center and median lanes, although less dramatic, followed close on the heels of the shoulder lane's trends, as evident in Figure 2(d). This shows that vehicle slowing in the shoulder lane spread quickly across the freeway.

And while the bottleneck remained active, shoulder lane slowing periodically re-occurred downstream of the on-ramp, although to lesser extent[4]. This important detail is exemplified in the figure next presented.

Figure 2(e) displays trip time ratios in the shoulder lane (as in Figure 2(c), but for a longer time) along with an oblique N-curve of counts departing the merge at MP 9.3 during this same period. The figure indicates that reductions in the latter tended to accompany rises in the former. The initial (and most pronounced) slowing that immediately followed the bottleneck's activation coincided with marked reductions in the flow, as highlighted with the smooth curve segments

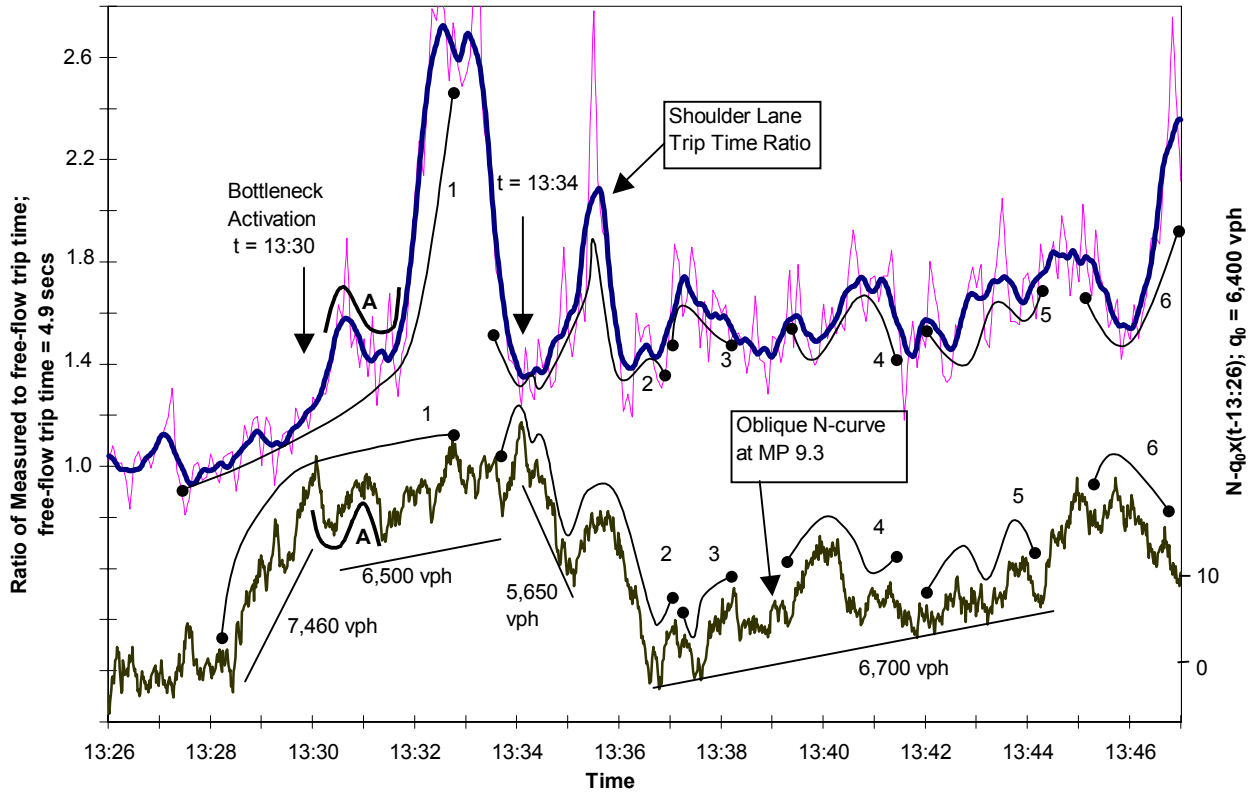


Figure 2(e)
Correlations between slowing and bottleneck discharge (March 15, 2001)

[4] Downstream slowing was even less pronounced in the center and median lanes after t = 13:34, such that the effects are scarcely evident in Figure 2(a).

labeled 1 in the figure. The latter dropped by nearly 1,000 vph, as highlighted with piece-wise linear interpolations of the N-curve.

Moreover, the slow-downs that occurred shortly after t = 13:34 accompanied further discharge reductions, as highlighted with the curve segments labeled 2. And the discharge flows eventually recovered (to an average rate of about 6,700 vph) once the shoulder lane trip time ratios consistently fell below 2.0. Yet even during this period, there seems to be evidence that these trip times were negatively correlated with bottleneck discharge; i.e., a convex trend in one was often accompanied by a concave trend in the other, as highlighted by the curve segments labeled 3 through 6. The slow-downs thus appear to have had deleterious effects on flows departing the merge.

These disruptive slow-downs, in turn, seem to have been linked to surges in the on-ramp's inflows. A sizable portion of vehicles from the on-ramp evidently moved into the freeway's center lane soon after having merged. Higher such lane-change rates were thus brought by each surge from the ramp. And, when these surges occurred in the presence of sufficiently high flows arriving from the freeway itself, the lane changing evidently became disruptive, such that the periodic slow-downs resulted. Evidence of this is provided in the following two figures.

Figure 2(f) displays shoulder lane trip time ratios along with an oblique N-curve of on-ramp vehicles as they passed MP 9.1 and merged. The latter shows that, on this day, on-ramp inflows surged to about 1,200 vph. And the figure shows clear correlation between these ramp surges and the shoulder lane trip times. Virtually every peak in the latter arose at or near the end times of the former and vertical arrows are included in Figure 2(f) to aid the reader in verifying this. And with few exceptions, each ramp surge that followed the bottleneck's activation corresponded to a rise in the trip times. Some linear approximations accompany the on-ramp's N-curve to highlight this.

Figure 2(f) also shows that the bottleneck's activation, the event that triggered a large measured reduction in flows departing the merge, was itself accompanied by a surge from the on-ramp. And by referring back to Figure 2(b), the reader can verify that this particular surge from the ramp was the first after the freeway traffic's arrival rate had risen to 7,320 vph. More will be said in the following section about the freeway and on-ramp flows that accompanied the bottleneck's daily activations.

For now, Figure 2(g) shows that the on-ramp surges tended to bring higher rates of lane changing (departing the shoulder lane) just downstream. An oblique N-curve of ramp inflows is

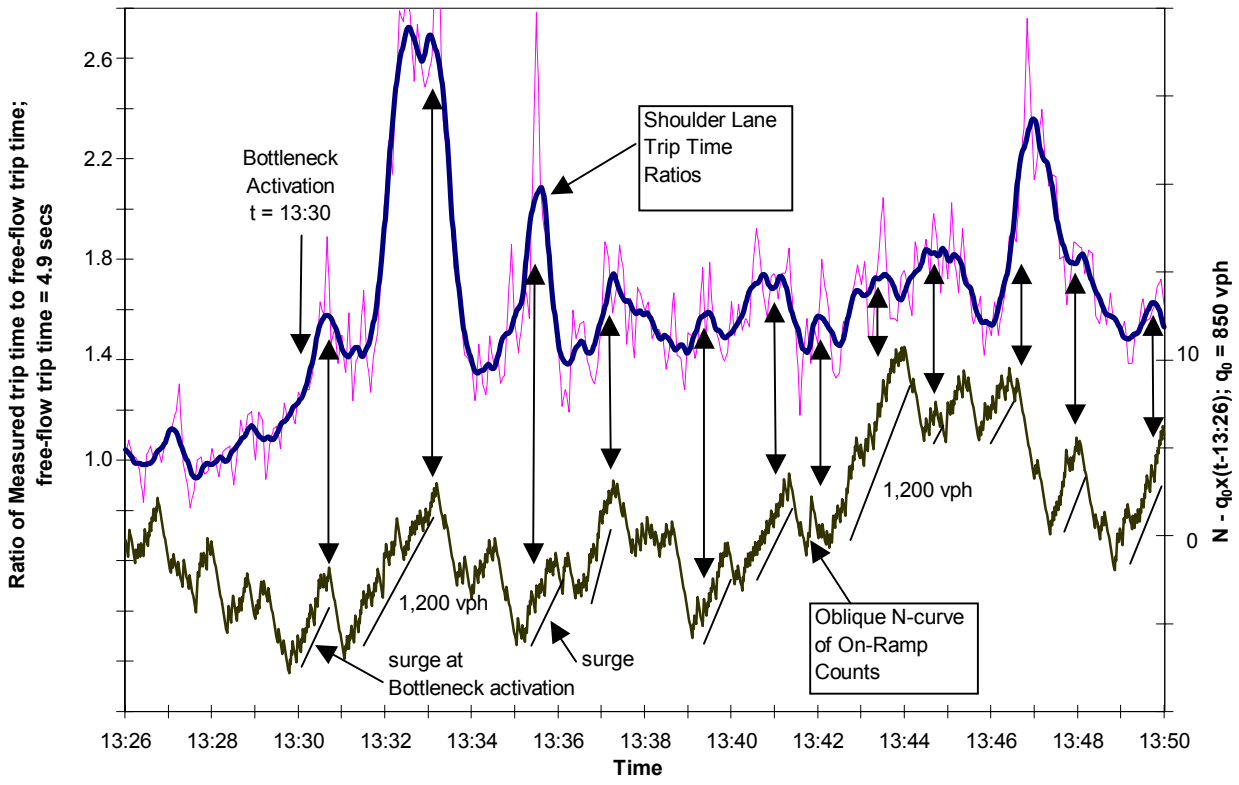


Figure 2(f)
Correlations between on-ramp surges and slowing downstream (March 15, 2001)

again included in this figure. Also shown is a curve displaying the net lane changing of interest; i.e., its slopes are the net egress rates from the shoulder lane between MPs 9.1 and 9.3. (Negative slopes on this curve mark periods with net movements into the shoulder lane). The lane-changing rates displayed here were estimated using N-curves of shoulder lane vehicle counts at MPs 9.1 and 9.3. These N-curves were constructed only during periods when trip times between these MPs did not appear to be varying markedly and the two periods shown in Figure 2(g) are instances of this.[5] By shifting the upstream N-curve for one such period forward by the average trip time between the MPs, the rates of (vertical) divergence between it and its downstream counterpart can be taken as approximations of the net lane-changing rates (Mauch and Cassidy, 2001). A thin line is used in Figure 2(g) to display these rates of divergence. The bold line shows their 10-sec moving averages.

[5] The reader can use figure 2(e) to confirm that the period from t = 13:37 to t = 13:44 was marked by relatively small peaks in the shoulder lane trip time ratios.

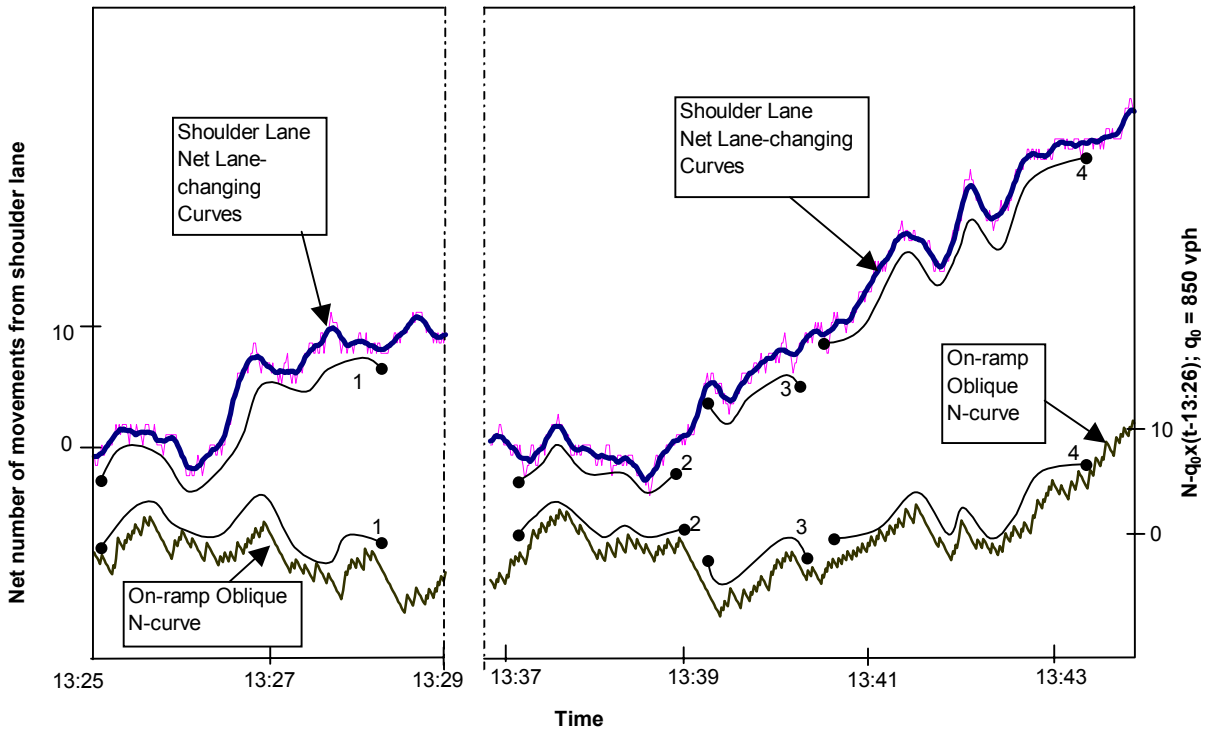


Figure 2(g)
Correlations between on-ramp inflows and lane changing downstream (March 15, 2001)

The on-ramp inflows shown in Figure 2(g) tend to display positive correlation with the net lane-changing rates. A convex (or a concave) trend in one is often accompanied by the same trend in the other. Smooth curve segments have been added to Figure 2(g) to aid the reader in seeing these coinciding trends.

Having identified the apparent effects on-ramp inflows can have on the merge area's traffic, it follows that metering these inflows can result in higher service rates at this freeway location. Evidence of this is provided in the following section.

4. OBSERVED BENEFITS OF RAMP METERING

In this section, data are displayed in ways that demonstrate on-ramp metering can be used to increase discharge flows once the bottleneck has activated. Data are also presented here to show that metering can postpone the bottleneck's activation and by so doing, further increase the merge area's service rates.

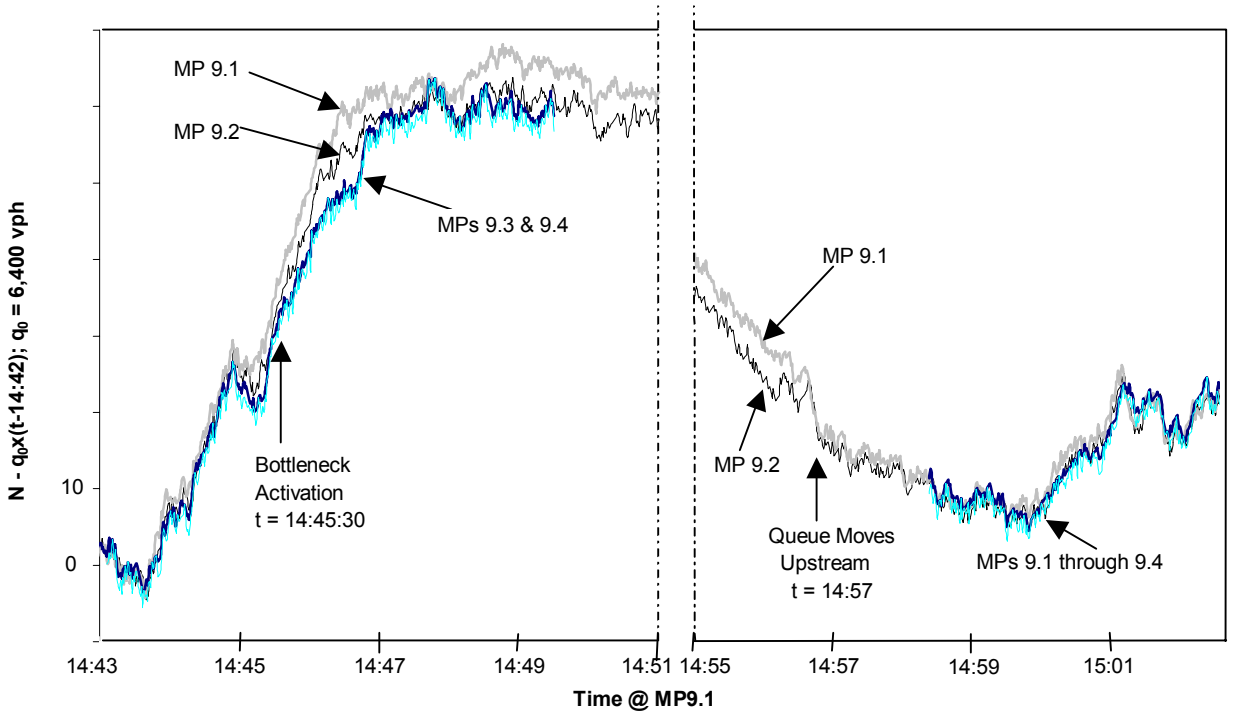


Figure 3(a)
Oblique N-curves at MPs 9.1 through 9.4 (February 20, 2001)

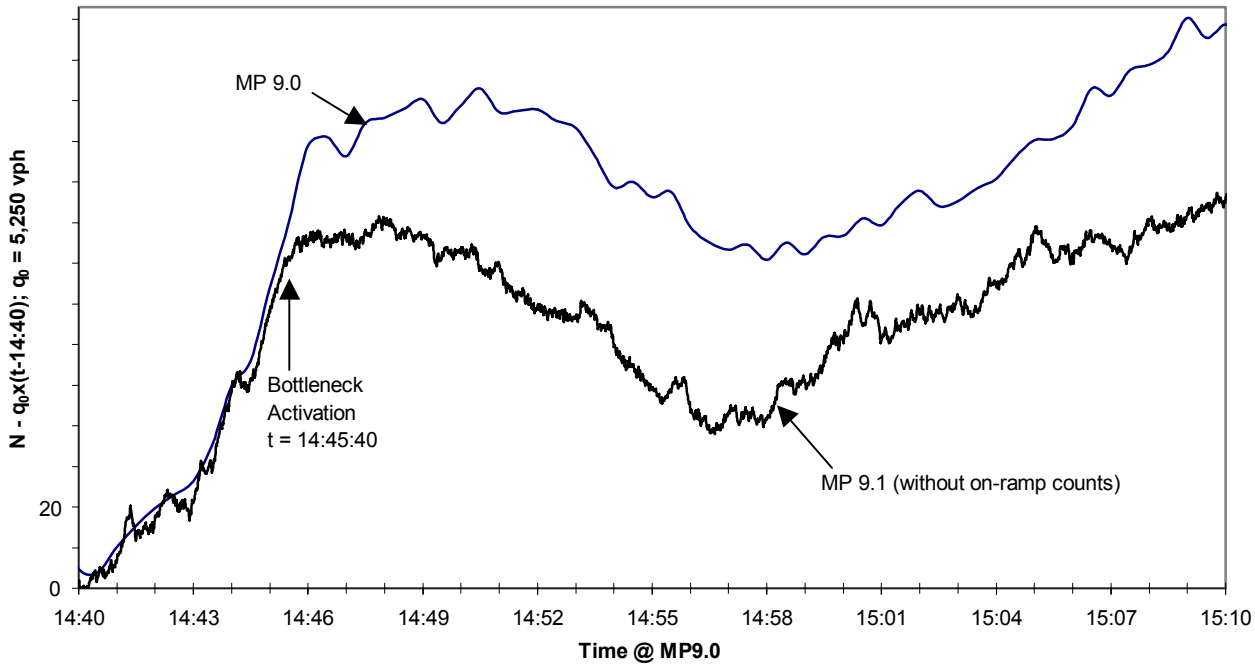


Figure 3(b)
Oblique N-curves of freeway traffic at MPs 9.0 and 9.1 without on-ramp counts (February 20, 2001)

Figures 3(a) and 3(b) are oblique N-curves from another observation day (February 20, 2001). The former displays portions of curves sufficient to verify that the bottleneck again activated between MPs 9.2 and 9.3, on this day at t = 14:45:30. It also shows that the head of the queue moved upstream by t = 14:57, leaving virtually no evidence of any delays downstream of MP 9.1 beyond this time. But Figure 3(b) demonstrates that queueing remained just upstream of MP 9.1 for some time thereafter.

Figure 3(c) presents the oblique N-curve of the on-ramp counts together with the shoulder lane's trip time ratios (just downstream) for a 30-min period during this rush. Prior to the bottleneck's activation (at t = 14:45:30), the peaks in the trip times closely correspond to the ending times of surges from the on-ramp; vertical arrows are included in the figure to highlight this.[6] The figure also shows that the bottleneck's activation was accompanied by a prolonged increase in ramp

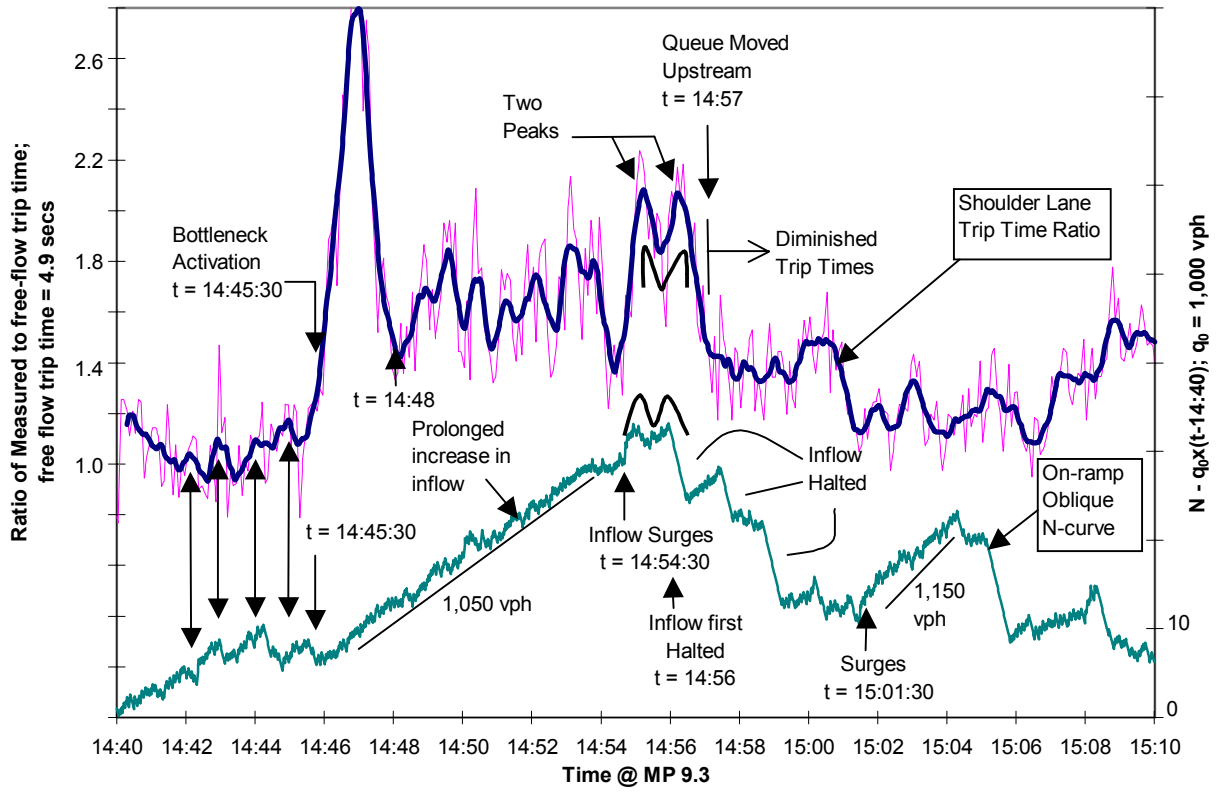


Figure 3(c)
Correlations between on-ramp surges and slowing downstream (February 20, 2001)

[6] Although the time scale used to display the curves in Figure 3(c) makes it difficult at times to verify by eye, similar correlations between trip times and ramp surges were also evident after the bottleneck's activation.

inflow and a marked rise in the trip times. The latter receded by t = 14:48, although ratios remained well above 1.0 for some minutes thereafter.

At t = 14:54:30, a rapid sequence of two inflow surges from the ramp caused trip times to rise (and display two pronounced peaks). But notably, these were followed by a sequence of three periods, each of about 30-secs in duration, whereby ramp inflows were completely halted. Consequently, the trip times diminished downstream; (these ratios were below those observed while the bottleneck was active for the day shown in the previous section). With this vehicle slowing all but eliminated, an eventual surge in ramp inflow (at t = 15:01:30) did not cause trip times to rise; they actually dropped slightly, such that ratios nearly returned to 1.0.

Figure 3(d) presents both the shoulder lane's trip time ratios and an oblique N-curve of counts departing the merge at MP 9.3. The figure shows that, at t = 14:47, the peak in trip times (brought by the bottleneck's activation) caused flows from the merge to drop to 6,130 vph. And the later rise in trip times at t = 14:54:30 caused discharge flows to drop even further to 5,680 vph.

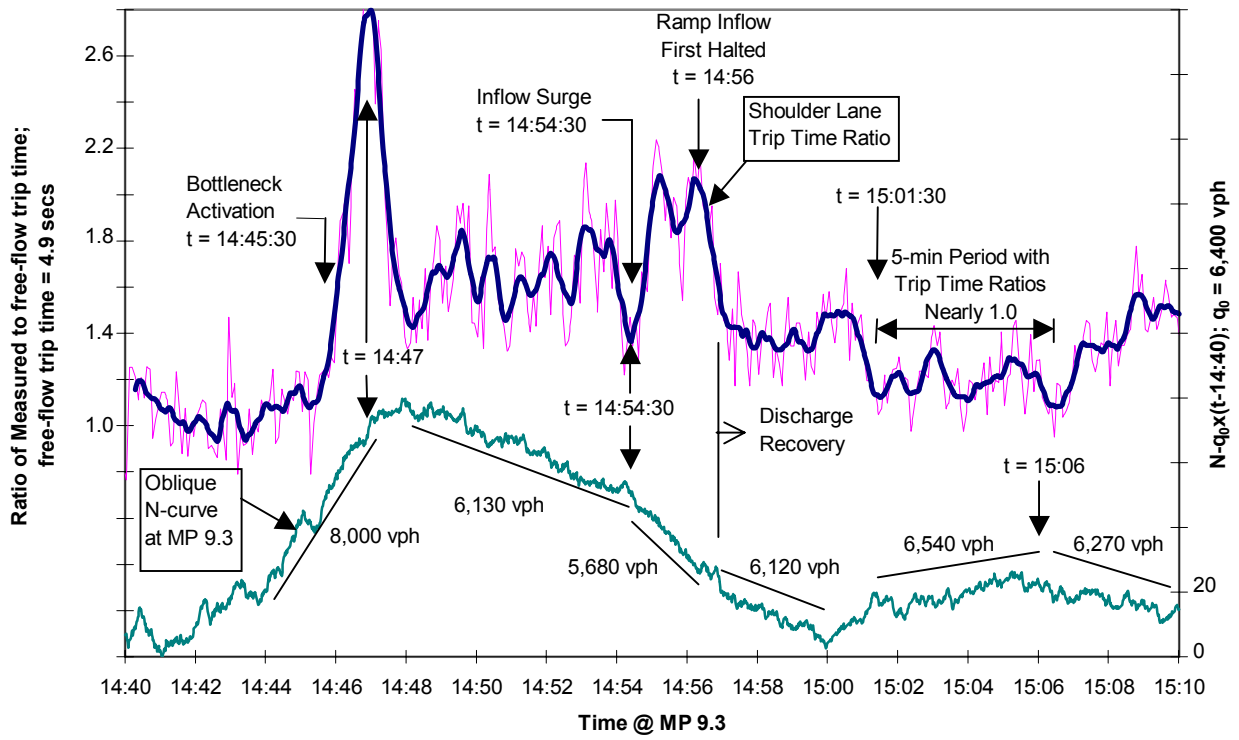


Figure 3(d)
Correlations between slowing and bottleneck discharge (February 20, 2001)

Discharge rates began to recover just after t = 14:56, the time when the ramp inflows were first halted. Moreover, the highest of these recovery discharge flows (6,540 vph) corresponded to the 5-min period when shoulder lane trip time ratios returned nearly to 1.0.

Also evident in Figure 3(d), the flow departing the merge just prior to the bottleneck's activation was 8,000 vph, a rate substantially higher than any of the queue discharge flows. Of course, this very high flow persisted for no more than a minute or two. But the data show that, prior to the bottleneck's activation, flows nearly as high can depart this merge for much longer periods.

An example of this is shown in Figure 3(e). Presented here is another oblique N-curve at MP 9.3, with this one displaying counts from earlier in the rush. It shows that the merge accommodated an average flow of 6,960 vph for a full 15 mins prior to bottleneck activation. This high rate was 10 percent greater than the average queue discharge that eventually ensued; (the latter was measured for an hour following this activation). This illustrates that higher service rates can be sustained at this merge by postponing the activation.

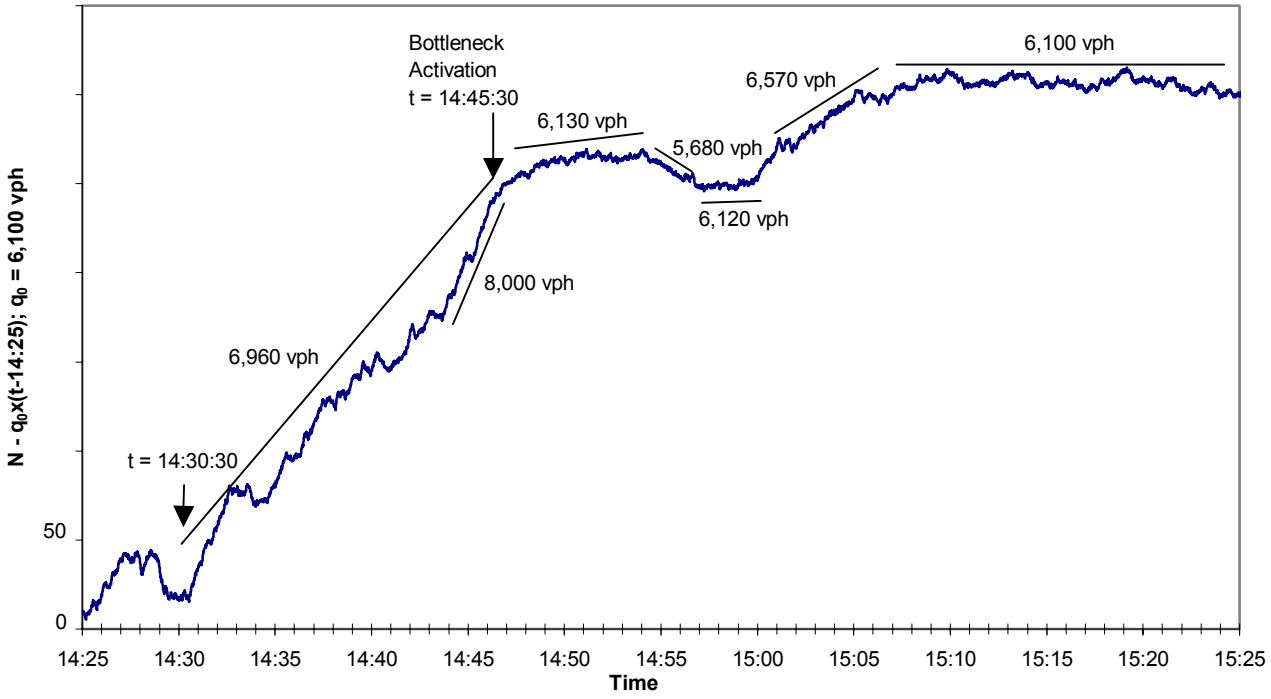


Figure 3(e)
Oblique N-curve at MP 9.3 (February 20, 2001)

Furthermore, the data in Figure 4 indicate that such postponements can be realized with suitable ramp metering schemes and that this can be advantageous. Displayed here are the on-ramp inflows and the corresponding flows departing the merge (at MP 9.3) measured at each day's activation. Their trend indicates lower values of the former can result in higher rates for the latter. Notably, the range of the measured departure flows is over 1,400 vph; it extends from 6,730 vph on November 8, 2000, a day when the on-ramp's meter was not in service (so inflows exhibited short-lived periodic surges of over 1,400 vph), to 8,050 vph on February 21, 2001, when the meter was more restrictive than on any of the other three days.

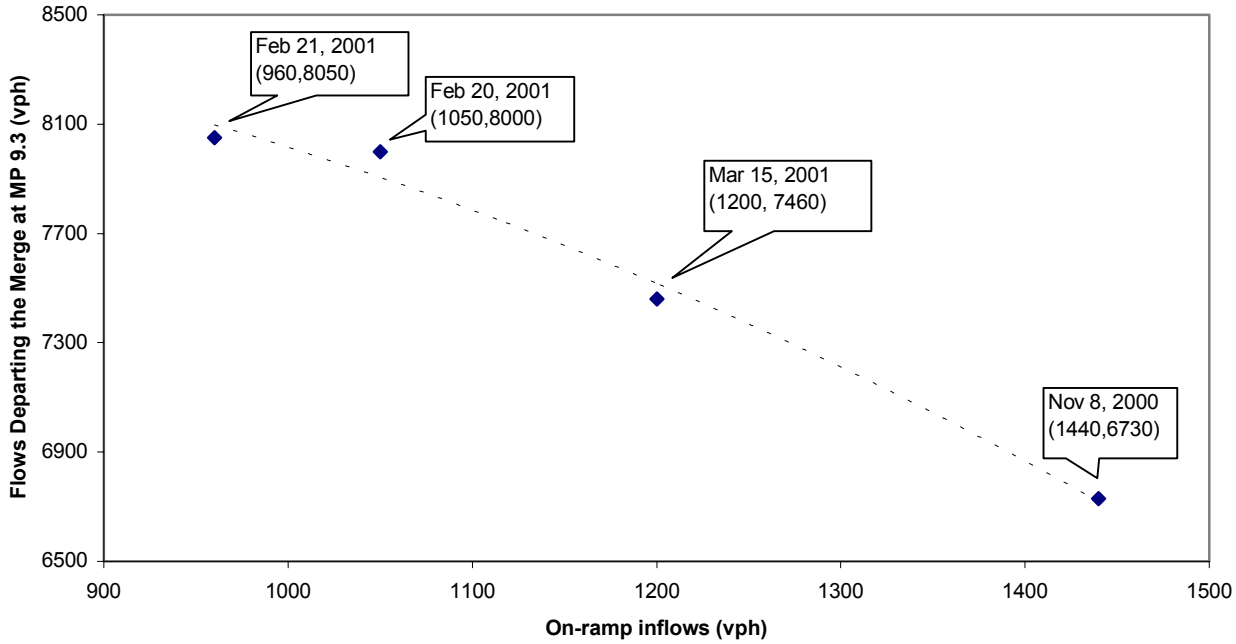


Figure 4
On-ramp and downstream flows measured just before the bottleneck activations on each observation day

Admittedly, the departure flows graphed in Figure 4 were not observed for prolonged periods. To the contrary, they each persisted for no more than 2 or 3 mins. But for the most part, each day's flows departing the merge gradually approached its rate shown in the graph. (The very high and persistent flow of 6,960 vph displayed in Figure 3(e) is a case in point). It follows that higher service rates can result from damping the on-ramp's inflows.

An ideal metering scheme would do this in ways that kept the flows departing the merge high without triggering bottleneck activation. Discussion on planned experiments for further testing of metering schemes is included in the final section; it follows next.

5. CONCLUSIONS

Metering freeway on-ramps is a widely used means of traffic management. It is well known that, under the right circumstances, metering can influence commuter travel decisions (e.g. route selection, trip scheduling, mode choice) to certain desirable ends, including system-wide delay reduction. And in almost any circumstance, metering can improve traffic conditions (e.g. facilitate higher travel speeds) on the freeway itself. This kind of improvement can be achieved merely by using the meter(s) to transfer queues from the freeway to the on-ramp(s). This can mean higher flows at locations where a queue would have otherwise resided to impede freeway traffic (MnDOT, 2001), particularly if the queue is kept from propagating past busy off-ramps and starving them of flows (Daganzo, 1996). But higher freeway flows well upstream of active bottlenecks are not evidence that metering increases bottleneck capacities.

The present study, however, has examined flows departing a merge bottleneck and has examined the potential for increasing the merge area's service rates by metering its upstream on-ramp. The findings indicate this might be realized by limiting in careful, deliberate ways the rate on-ramp vehicles join the freeway's traffic stream. This was made particularly clear in Figure 4. It follows that such metering schemes can reduce the delay collectively incurred by vehicles, even absent any changes in commuter travel patterns in response to the metering.

These findings, while very important in their own right, are not the final word on the subject, however. Future experiments at this site will further test metering rates suitable for postponing (or perhaps even eliminating) the bottleneck's activation. Also to be tested further is the potential for raising discharge flows by metering more restrictively whenever shoulder-lane slowing is detected subsequent to the activation. By temporarily halting the on-ramp's inflow at the first sign of such slowing, it may even be possible to restore freely flowing conditions (and higher service rates) at the merge and this too will be tested.

But even if these strategies prove successful for the present site, they would be appropriate only for bottlenecks that display similar traffic features; i.e., bottlenecks triggered by on-ramp vehicles that change lanes soon after merging. At present, the extent to which other merge bottlenecks exhibit this crucial detail is unknown. It is notable that previous study of merge junctions on freeways near Toronto, Canada reported that bottlenecks activated some distances downstream of their on-ramps (Cassidy and Bertini, 1999). This could be explained by lane-

changing behavior consistent with the present findings. Regrettably, however, the Toronto data came from loop detectors with rather coarse resolutions that precluded confirmation of the bottlenecks' cause(s).

To complicate matters further, an even earlier study of a merge on another California freeway (Banks, 1990) reported a bottleneck's location diametric to our present findings; i.e., rather than observing the bottleneck's initial activation in the shoulder lane just downstream of the merge, this earlier study reported that queues first formed in the median lane at a location well in advance of the on-ramp. However, this previous study did not provide evidence of this claim; there was nothing comparable to the queueing diagrams used in the present manuscript, for example, to verify the bottleneck's location. It is therefore not clear if the earlier (California) study points to real variations in traffic features across freeway merge bottlenecks or if subtle details like those reported in the present manuscript went undetected in the earlier work.

This uncertainty underscores the need to study additional freeway merge bottlenecks using high-resolution data treatment methods like the ones described here. Such studies are currently under way.

REFERENCES

Banks J.H. (1990) Flow processes at a freeway bottleneck. *Transpn Res. Rec*. 1287, 20-28.

Cassidy M.J. and Windover J.R. (1995) Methodology for assessing dynamics of freeway traffic flow. *Transpn Res. Rec.* 1484, 73-79.

Cassidy M.J. and Bertini R.L. (1999) Some traffic features of freeway bottlenecks. *Tranpn Res.* 32B, 49-59.

Daganzo C.F. (1996) The nature of freeway gridlock and how to prevent it. *Proc. Int. Symp on Transportation and Traffic Theory.* (J.B. Lesort, ed), pp 629-646, Elsevier, Oxford, U.K.

Daganzo C.F. (1997) *Fundamentals of transportation and traffic operations*. Elsevier, New York, pp 259-261.

Mauch M and Cassidy M.J. (2001) Freeway traffic oscillations: observations and predictions. *Submitted for publication.*

Minnesota Department of Transportation, MnDOT (2001) Twin cities ramp meter evaluation, Executive Summary. Cambridge Systematics, Inc. *Available on the world wide web at* http://www.dot.state.mn.us/rampmeterstudy